SANDIA REPORT
SAND88—2416 • UC—605
Unlimited Release
Printed May 1989

Network Architecture Functional Description and Design

Version 2.0

Len Stans, Mike Bencoe, Doug Brown
Sue Kelly, Lyndon Pierson, Charlene Schaldach

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550
for the United States Department of Energy
under Contract DE-AC04-76DP00789

SF2900Q(8-81)

When printing a copy of any digitized SAND Report, you are required to update the markings to current standards.

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, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents 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
PO Box 62
Oak Ridge, TN 37831

Prices available from (615) 576-8401, FTS 626-8401

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

NTIS price codes
Printed copy: A05
Microfiche copy: A01

SAND 88-2416
C-32

Network Architecture

Functional Description and Design

Version 2.0

May 25, 1989

Len Stans, Project Leader
Mike Bencoe, (Original Team Member)
Doug Brown
Sue Kelly
Lyndon Pierson
Charlene Schaldach

Abstract

This report provides a top level functional description and design for the development and implementation of the central network to support the next generation of SNL, Albuquerque supercomputer in a UNIX® environment. It describes the network functions and provides an architecture and topology.

Acknowledgements

The Network Architecture Project included the following personnel:

Len Stans, 2640A, Project Leader
Mike Bencoe, 2645 (Original Team Member)
Doug Brown, 2645
Sue Kelly, 2645
Lyndon Pierson, 2647
Charlene Schaldach, 2643

This project had the participation of a large percentage of 2600 personnel and management, who reviewed many draft documents and attended numerous design review meetings. Their ideas and suggestions were a significant contribution to this project.

Larry Buxton, Supervisor 2643, carefully reviewed many draft copies of the report for technical consistency, readability, and grammar.

Parts of the glossary were extracted from a glossary of computing terms published in the Sandia Computing Newsletter, compiled by Gary Shepherd and other staff on Division 2614 and from the "Scientific Workstation Project Final Report" (SAND88-1179).

This document contains references to the trademarks of the following corporations.

Corporation	Referenced Trademark
AT&T Bell Laboratories	UNIX®
Digital Equipment Corporation	DECnet, VAX, VMS
International Business Machines Corporation	PROFS, SNA
Masstor Systems Corporation	MASSNET
Network Systems Corporation	HYPERchannel, NETEX, DX, Network Systems
Xerox Corporation	Ethernet

Table of Contents

		Page
1.0	Management Summary	5
2.0	Introduction	7
3.0	Constraints and Assumptions.	8
4.0	Future Role of the 2600-Supported Computing Resources	10
5.0	Functional Requirements of Network Services	12
	5.1 User Authentication	14
	5.2 Terminal Services	15
	5.3 Network Time Service	18
	5.4 Network File Transport	19
	5.5 Gateways/Bridges	20
	5.6 Batch Job Scheduling	23
	5.7 Network File Storage Service	24
	5.8 Task-to-Task Communication	26
	5.9 Output Processing	27
	5.10 Network Monitoring and Status Reporting	27
	5.11 Network Mail & Messaging Facility for the SSN/PON	27
	5.12 Accounting/Billing Service	29
6.0	Network Architecture Topology	30
	6.1 Hardware.	33
	6.2 Standard Host Interface	33
7.0	Inter-Network Relationship with; SNL, Livermore; TTR; and other Sites	39
8.0	Protocol Support for the SSN/PON	41
9.0	Network Security	45
	9.1 Network Security Policies	45
	9.2 Protocol Support for Security	47
	9.3 Encryption as a Tool for Network Security	47
	9.4 Separate Physical Subnetworks	47
	9.5 Security Filters	48
	9.6 Security Partitions	49
	9.7 Network Security Design	49
10.0	Network Performance	50
	10.1 Data Transfer Rates within the SSN/PON	50
	10.2 Data Rate Requirements for Supercomputer-to-Workstation Communications	50

11.0 Network Operations/Management & Maintenance 53

12.0 Network Conventions/Standards 54

13.0 Migration Considerations 55

14.0 Network Reliability and Maintainability 57

15.0 Design Document Revision and Control 58

Appendix A 59

Appendix B 67

1.0 Management Summary

This document primarily describes the proposed network that will support SNL, Albuquerque's new Supercomputer-90 system in a secure UNIX environment through the mid-1990's time period. The network performance specifications are the initial values to be implemented to support the Supercomputer-90 system. It is anticipated that the architecture and implementation designs will continue to be enhanced over the next 3-5 years to provide higher network bandwidths and throughputs, as the hardware and protocol technologies evolve.

In viewing the Central Computing Network (CCN), it is important to recognize that, conceptually, the current CCN has become a subnet in a larger network consisting of user Local Area Networks (LANs) and peer networks. The CCN currently supports two major functions. One function provides selected services directly to the distributed users, such as file storage and output processing. The other function integrates CCN services with selected CCN workers into a single computational environment. In general, the bandwidth and protocol demands placed on the CCN to support these two functions can be substantially different.

The development of the future network will be approached from the perspective that 2600 computing will continue to evolve from the current position of being the "central site" toward a position of being a unique subnet on a larger network (herein called the SNL Internet). The SNL Internet will ultimately consist of a number of interconnected independent nodes and LANs, each providing and utilizing network services and functions.

The requirements for the architecture of the network are influenced by a number of factors: the future Supercomputer-90 user requirements; the availability of high performance workstations; the goal to utilize industry standard protocols and operating systems; and finally the general requirement to provide improved throughput, security, reliability, and functionality.

It is envisioned that the new network will be an integral part of a UNIX based computational environment consisting of industry standard and vendor supported hardware and software components. TCP/IP, NETEX, and DECnet (DECnet will be used primarily on distributed networks) will be the initial protocols, combined with a high bandwidth network backbone based on Network Systems Corporation (NSC) network hardware.

Security constraints and system functionality dictate the construction of separate physical computer networks for the dissimilar computing environments. This architecture specifies that a secure network (per DOE Order 5637.1), an administrative network (per DOE Order 1360.2), and a protected, open network (per prudent business practices), be implemented. As used in this document, the adjective secure is defined as indicating an environment in which all users are Q-cleared. The adjective open is defined as indicating an environment in which the users may not be Q-cleared. Protected means that control mechanisms have been developed and implemented to safeguard resources. Communication between the three networks would be by way of filters (e.g., dual--ported disks, encryption, administrative procedures) that would provide controlled, audited, and managed communication paths between the separate networks/partitions.

The new networks being described would be distinct from the current CCN. A file transfer facility between the existing CCN and the new secure network is proposed. The secure distributed VAX network will initially have connectivity to both the CCN and the new secure network. The CCN connectivity will be discontinued when all the needed functionality is available on the new secure network.

The proposed network architecture is designed primarily to support a supercomputer with a UNIX-based operating system. The network services are specified to provide the performance and capacity necessary for user support in such a supercomputer environment. A secondary, but important, goal is that many of the specified network services will be distinct, modular functions that will be available to other interconnected LANs, almost as if they were part of those LANs. By using a homogeneous UNIX-based networking environment, the details of the network topology should be transparent to the user.

Chapters 1-4 of this document include the management summary, the introduction, the constraints and assumptions that guided the design, and the future role of the 2600-supported resources. Beginning with Chapter 5, the functional description and design are covered.

Chapter 5 describes the functionality of the network. The individual functions comprising that functionality are:

a. User Authentication
b. Terminal Services
c. Network Time Service
d. Network File Transport
e. Gateways, Bridges, and Routers
f. Batch Job Scheduling
g. Network File Storage Service
h. Task-to-Task Communication
i. Output Processing
j. Network Monitoring
k. Network Mail
l. Accounting/Billing Service

Chapter 6 presents the topology of the network. It describes an NSC "DX technology" HYPERchannel backbone plus additional subnetworks for terminal, security and operational management traffic. Chapters 7 through 15 provide additional details of the architecture including inter-network relationships, protocols, security, performance, management and operations, migration, implementation, reliability, and document revision.

The migration plan to implement this architecture will be maintained as a separate document where the functions described in this document will be converted into designs, which will in turn be converted to implementation tasks. Each task is to be initiated as a distinct activity which will result in a production implementation. The production implementation for the secure network, which will support the Supercomputer-90 system has the highest priority. Implementation of the open network will be accomplished as the open computing requirements are better established and as resources become available.

2.0 Introduction

The Statement of Work for the Network Architecture Project (dated July 21, 1987) defined the design objectives for the network to be:

a. Support the Supercomputer-90 in a secure UNIX environment.

b. Support the direct and/or indirect connection of workstations.

c. Provide a network architecture that is consistent with the Directorate 2600 strategic plan.

d. Provide connectivity for the administrative and office automation subnetworks.

The requirements document (dated December 4, 1987) specified:

a. The network implemented by the new architecture must operate in parallel with the current network to minimize the disruption to network users.

b. There must be multiple network interfaces to critical worker and support nodes in order to minimize single points of failure.

c. Network security must be improved through the implementation of separate networks based on security partitions (open, secure, and administrative information), and the implementation of an overall network security architecture.

d. The network must support a mixture of "worker" and "service" machines.

e. The network must be flexible and expandable to support new network services.

f. Hardware and software components must be industry-standard and supported by at least one stable vendor.

g. The scientific networks must support UNIX operating systems.

h. The network must support a common user environment between SNL, Albuquerque and SNL, Livermore.

This document provides the functional design that was specified in the Statement of Work (SOW). The implementation plan will be derived from the individual task SOWs.

3.0 Constraints and Assumptions

The major assumptions made in the development of the scientific network's functional design and requirements are the installation of SNL, Albuquerque's Supercomputer-90 system and the implementation of a secure UNIX environment. Other assumptions and constraints that influenced the functional requirements and design are:

1. The new network will co-exist, as a new and separate network with the current network.

2. A bidirectional file transfer facility will be developed to pass data between current network applications and the new network applications.

3. The future network topology will consist of three separate information partitions implemented as three physically separate networks: Secure Supercomputer Network (SSN), Protected Open Network (PON), and ADministrative NETwork (ADNET). Each of the physical networks may consist of multiple subnetworks to satisfy the required functionality, redundancy, capacity, and security.

4. The emphasis of this design is on the implementation of the SSN for the support of the Supercomputer-90 system. Therefore, most resources will be directed at the SSN. It is intended that the PON be implemented utilizing the same basic developments used in implementing the SSN, with appropriate adjustments for specific hardware differences when the resources become available. Therefore, any completion dates indicated herein for the PON are tentative.

5. The new network implementation will support multiple protocols. At a minimum, initial support for protocols will include TCP/IP, NETEX, and DECnet. The general strategy for protocol support is to structure the new network in such a manner as to provide for the evolution to new standard protocols as they become available and to be able to employ special purpose protocols where performance or functional requirements dictate. The old network will continue to support a combination of MASSNET, NETEX, and DECnet.

6. The new networking hardware for the machine-to-machine backbone will consist of components from Network Systems Corporation (NSC). The design will provide for the inclusion of other networking hardware provided by other vendors.

7. The worker nodes on the new network backbone will execute UNIX-based operating systems. The server nodes must support this UNIX environment.

8. Supported, distributed networks <u>will</u> <u>not</u> be required to execute UNIX-based operating systems. However, it is expected that, over time, UNIX-based operating systems will become the dominant operating systems for distributed users.

9. The implementation of network functions will be evolutionary. Further, the functions described in this document are a combination of required and highly desirable functions. Therefore, it is possible that some of the desired functions described in the

document may not be implemented or their implementation may be delayed due to resource limitations. The specific implementation details for each of the functions will be discussed in the individual implementation plans for each of the network nodes or subsystems.

10. Values provided for transmission rates, storage densities, etc., are the targets for the initial support of the Supercomputer-90. It is anticipated that larger values will be required to satisfy expanded and new requirements over the next three to five years. The initial network rates discussed are based on the minimum performance needed to support the proposed services on the Supercomputer-90, the experience of the Workstation and the Graphics projects, and the projected availability of industry supported hardware and software.

    In the implementation phase, the Network Architecture project recommends the implementation of the maximum hardware capabilities for the funds available. However, due to the dynamic nature of the industry, it is not possible to explicitly specify the minimum achievable performance for all of the hardware components. For the backbone network, we believe a minimum network bandwidth is 100 Mb/s and for the subnetworks a minimum is 10 Mb/s. Through the replication of these minimum networks, higher aggregate bandwidths can be achieved. If hardware to achieve these minimum bandwidths can not be obtained, then the affected tasks should be put on hold until the required networking capabilities are available.

11. The Network Architecture project team is recommending that a single design approach be pursued, as opposed to several alternate designs, based on the requirement to maintain and migrate from the current network, the available vendors' networking products, the short (18-24 month) implementation schedule, and the available resources. Contingencies will be developed as part of the implementation plans and designs. Additional performance information is required from the network hardware vendors and the Prototype Secure Network project before specific contingency plans can be fully developed.

12. The migration strategy for evolving from the current network to the future network will have the current CCN continue to function in parallel with the new network, with the implementation of a <u>file</u> migration facility between the two networks. We will not attempt to implement the existing network functions such as MASS, AFT, etc. into the new network because the new network is based on a UNIX environment that has a set of equivalent functions. We do not believe the two sets of functions are sufficiently compatible that they would allow any reasonable migration of the actual functions from the old environment to the new environment.

4.0 Future Role of 2600-Supported Computing Resources

It is envisioned that over the next several years the computing environment will continue to shift more and more toward distributed network computing. Just as much of the small and medium-scale computing has shifted toward distributed VAXes during the past eight years, the availability of low-cost high-performance workstations will drive computing at Sandia toward the user's desk. We project that most of the users at SNL will be able to perform all of their computations on their own workstations, which will have speeds in excess of 10 MIPS. Use of the supercomputers in the central site LAN will be limited to that class of problems that require very large memories and/or computational power that will not be available on a high-end workstation or departmental compute server.

Departmental computers will become the hubs of Local Area Networks to which high-end workstations are attached. Their computational role will decrease, so that they probably will become more like network routers and gateways to a larger SNL secure (or protected) network. The departmental computers are also likely to become file servers for the attached workstations in order to facilitate file sharing, backup, and media control by a departmental system management group.

This means that the role of the central site as a computational node may decrease, so that the supercomputers will serve a relatively small (but still important) user community. The 2600 networks will thus become another subnetwork along with many other subnetworks that are, in fact, a larger network that we shall call the SNL Internetwork. Because of the change in emphasis, we propose that the term CCN be abandoned for the new network. Suggested scientific network names are Secure Supercomputing Network (SSN) and Protected, but Open, computing Network (PON). If these network topology recommendations are implemented, the administrative network (ADNET) will exist as yet another separate network. We will provide for a controlled and managed communications connection between the ADNET and the SSN and/or PON.

****

The purpose of having two scientific networks, the SSN and the PON, is to limit the SSN to users with Q-clearances. Users of the PON may or may not have Q-clearances. This document follows the SNL, Albuquerque tradition of using the adjectives secure to mean a computing environment in which direct or indirect access to computer resources is restricted to Q-cleared personnel. Open means potential access by non-cleared individuals. The adjective protected is used to mean that prudent business practices (e.g., password control, dial-back mechanisms, etc.) are employed in the development and operation of the system.

****

The primary functions of the Network Architecture will be to:

1. Provide state-of-the art supercomputing to those users who need to run very "large" problems.
2. Provide archival file storage and a central repository for shared data.
3. Provide specialized high-cost output services, e.g., microfiche, specialized graphics.
4. Provide network routing and gateway interconnections within the SNL Internetwork.

5. Provide a communications connection for intersite networks to/from LANL, KDC, etc. This connection is intended to maintain the current level of service between these sites in the new network.
6. Provide for central network management and support to aid in the operation of local LANs and workstations through central network diagnostic utilities, and distributed system software support.

Some of the generic kinds of services provided by the SSN and/or PON are specifically directed at providing support to and for the scientific computers and supercomputers, since historically it has been our experience that state-of-the-art supercomputers come from the vendors with minimum software and hardware environments needed for a complete user environment.

Most of these services are also intended for general use by users throughout the SNL Internetwork. The services being provided to the internetwork users, generally, will not have the same performance requirements as those that are provided in support of the supercomputers. However, there is a considerable overlap between the services needed by distributed computer and workstation users and the supercomputer users. By interconnecting the distributed networks with the SSN/PON, we can provide access to the SSN/PON services from distributed systems and workstations, and we can also allow users on different distributed computing networks to communicate with each other. Gateways and/or routers will perform the protocol conversions and security checks that allow the secure interconnection of many subnetworks that are in the same partition.

5.0 Functional Requirements of Network Services

The functional requirements for the environment provide the foundation upon which the network architecture and topology will be built. This section describes the general functional requirements that the network must support. The specific architecture and topology will be discussed in Chapter 6.0.

The network functions are to consist of a set of physically independent functions that operate in a logically integrated manner. The purposes of a majority of the services are two-fold: 1) to provide indirect user support to the supercomputer(s) users and 2) to provide independent support directly to distributed network users. The performance requirements of each of the services will, in general, be driven by the more severe support requirements of the supercomputer. However, if the services are modularly designed, they can be utilized by other network users, nodes, and LANs.

This chapter identifies the specific services or classes of service which we would expect to be available in the new networking environments (SSN and PON) and which would be supported by the new networks. The functionality of each of these services is being described here without attempting to provide the implementation design, constraints, etc.

a. User Authentication
b. Terminal Services
c. Network Time Service
d. Network File Transport
e. Gateways/Bridges
f. Batch Job Scheduling
g. Network File Storage Service
h. Task-to-Task Communication
i. Output Processing
j. Network Monitoring and Status Reporting
k. Network Mail and Messaging Facility
l. Billing Service Input (input to central billing)

There are two other generic functions that were identified: a Magnetic Tape Input/Output Service, and a Data Conversion Service. They have not been included in this list since we could not identify a specific requirement for these services in the new environment.

Unless a constraint or limitation is stated explicitly, each function potentially may be installed in any of the three networks. For example, gateways might be a legitimate function in secure, protected, or administrative networks, to support user requirements.
Each potential network service has the following common functional requirements, without regard to the networks in which they are installed:

1. Each service must be able to execute in a UNIX environment with little or no customization required. This implies that a network service which is heavily dependent on the internals of a proprietary operating system (e.g., IBM/MVS, VAX/VMS, etc.) would not be a choice for implementation, unless extenuating circumstances mandated it.

2. Each service must identify its robustness and reliability requirements. The design must describe how the requirements will be satisfied.

3. Each service must include gathering (and possibly processing) of performance data and data on internal failures and retries. While Section 11 defines the need for an overall performance and statistics gathering function, there will always be a need for service-specific statistics.

4. Each service must include an audit feature. This feature should audit both access and usage. If this service is implemented in vendor-supplied software, preference should be given to an audit facility which is easily tailored and does not require source changes in the product to modify it to site specifications.

5. Each service must include an accounting/resource utilization facility. This facility should create usage records and have a tailorable billing feature, particularly if charge-back is desired. For a vendor-supplied implementation, the accounting facility should satisfy the tailoring criterion stated previously for the audit feature.

The issues or functional requirements for implementing most of the basic services in a secure network are discussed in the following subsections for the individual services.

5.1 User Authentication

A user authentication service is necessary in any environment, whether it be imposed by a DOE order or prudent business practices. The definitions for authentication, authorization, and validation are provided for commonality of use.

Definition: Authentication - The act of verifying the claimed identity of an individual, station, or originator [of a computing process]. (DOE Order 5637.1)

Definition: Authorization - 1). The privilege granted by an appropriate official to access a particular classification level or category of information within an ADP system (DOE Order 5637.1). 2). The process of granting services provided by computing resources on the basis of the identity of a person (determined by authentication) and the recorded decision of an appropriate official to grant access privilege.

Definition: Validation - refers to both authentication and authorization (definition 2 above) implemented together as a single process or procedure.

The functionalities needed from a user authentication service are:

1. To establish user identity by means of passwords or other appropriate means and to provide non-forgeable tokens of this user authentication to the various network services as a basis for their granting (authorizing) specific services.

   The tokens of user authentication may represent authentication by more than one authenticator (some conveying greater trust). For example, a remote job entry service may trust a token of authentication generated by an authenticator on the worker machine. However, a network file service may require a token of authentication generated by a dedicated (more trusted) central authenticator to base its authorization decisions upon. This also allows the use of multiple authenticators (at each trust level) for reliability via redundancy.

2. To provide a capability for electronic signatures, using the non-forgeable tokens of user authentication. This is a desirable feature but probably will not be available in the short term, as the state-of-the-art has not yet advanced to the point where this is practical, at least in the specific form needed.

3. To use industry standard authenticators as soon as industry standards evolve and are embraced by DOE.

5.2 Terminal Services

Terminal services are those services in which a user interacts with a resource and expects a nearly instantaneous response time. Common terminal services include accepting logins and commands from a user, and returning responses to the command from the resource that the user is interacting with. Traditionally, the resource that a user issues commands to has been a single computer and its associated peripherals. The user has communicated with the computer through a dumb terminal and modems connected via the RS232 interface standard. Responses communicated to the dumb terminal typically range from a few characters to a screenfull of simple characters. A few terminals allowed full screen high resolution graphics images to be displayed.

More recently, PCs and workstations are being used to emulate the traditional "dumb terminal". The emulation functions exactly like the original terminal, with the added capabilities of being able to emulate more than one terminal device and to upload or download character or binary files from or to local storage.

With the introduction of workstations and network services (like DNTS), terminal services merge with computer-to-computer communications. The resources may become a network, including compute servers, I/O servers, and file servers, rather than a single computer and its peripherals. Although the commands issued to a resource are expected to remain a few characters long and the command could result in the movement of files or complex graphics information, terminal response should not be confused with file transfer service.

5.2.1 Implementation Considerations

Currently, there are two views on how terminal services should be implemented. There are circuit switched terminal services, (e.g., the traditional Terminal Switching Network services), and packet switched terminal services (e.g., such as those available through TCP/Telnet or DECnet/Set Host). There are advantages and disadvantages to each of these types of service. Our current circuit switching technology readily permits authentication on the basis of user id and terminal location, provides good directional control of connection initiation (terminal-to-host only), provides good need-to-know separation, uses Sandia's existing wire plant, supports high aggregate bandwidth, but provides less than 56 kb/s throughput per connection and does not readily allow multiple simultaneous connections (windows) to various hosts from a single terminal. Current packet switched (virtual terminal) technology can give users shared access to line rates of greater than 56 kb/s, readily provides multiple simultaneous connections (windows) from a single terminal, but provides poor need-to-know separation, and does not allow for authentication on the basis of terminal location.

Circuit switched technology (for some of the reasons cited above) permits access to multiple information partitions from a single (secured) terminal location, while packet switched technology currently does not. This creates the need to prevent simultaneous access to the TSN and to LANs. This is because LANs currently must exist in only one information partition and LAN terminals simultaneously connected to the TSN may form an unauthorized, unauditable bridge between information partitions. These security issues need to be resolved to provide the workstation user flexibility in terminal services. To

achieve flexibility, it will be important that workstation users have access to both circuit switched and packet switched facilities. The interactive path between a user and his computing resources will likely be composed of a combination of both technologies. It appears that over the next 2-3 years, with the technologies being devised by PBX manufacturers, computer communications vendors, and telephone companies, that the feature differences between circuit switched and packet switched services will blur.

5.2.2 Requirements for Terminal Services

It is anticipated that the traditional terminals will be augmented and/or replaced by PCs and workstations that are connected via LANs into distributed computing networks. The terminals, PCs, and workstations may have network virtual terminal access to the SSN/PON workers via gateways or routers. The separation and management of terminal services and information between secure and non-secure systems will be accomplished through connection and password management.

The Administrative Network (ADNET) will require both synchronous as well as asynchronous terminal support. That will be provided as part of the ADNET implementation effort.

The requirements for terminal services in the SSN/PON will be:

1. Short response times (as opposed to high throughputs) for a range of services from traditional terminal services to complex graphics display.

2. Access controls and other security features including:
   a) Authentication and authorization of access to resources.
   b) Auditing of user location and identification of requests for resource access.
   c) Auditing of network topology and changes to topology.
   d) Controlling of access to the switch or network.

3. Reasonably transparent service to the users by presenting a common terminal path setup for a wide variety of resources.

4. Support for multiple simultaneous terminal sessions without significantly degrading services.

5. Support for multiple sessions to a single resource.

6. Ubiquitous service - available everywhere, and easy to move, ability to add or remove terminals or workstations.

7. Simple to manage and maintain, requiring no increase in maintenance staff. This will be accomplished by having appropriate, commercially available, support and diagnostic tools.

8. Ability of terminals to transmit ASCII characters to initiate connections or path setups.

5.2.3 High Speed Graphics/Workstation Terminal Services

High speed graphics terminals and workstations will need to be circuit switched either directly or indirectly to a terminal server for supercomputer access. The workstations or PCs may use a network virtual terminal protocol such as TELNET to access the SSN/PON workers from the terminal server. The primary requirements of a high speed graphics/workstation terminal service are high bandwidth and good accessibility.

Bandwidths of at least 1 to 10 Mb/s will be required to support some high-end workstations and graphics terminal connections. Other technologies such as FDDI may be required to support multiple workstations at these higher data rates.

The new digital PBX will support large numbers of 1 Mb/s Ethernet connections (up to 3000) within the next 5 years and between 500 and 1000 connections by mid-1990. This may suffice for connecting many of the workstations to remote services, but may not be sufficient for connecting some high-end workstations to the SSN. At present, only a small percentage of the total number of terminals must access this high speed service from almost any location in the SNL Tech Areas, much as circuit switched terminal service is now.

5.3 Network Time Service

The SNL Internetwork requires a consistent clock. Diagnosing errors involving multiple network nodes may depend heavily on identifying the sequence of events that surround the problem. Tracking a possible security breach also requires time stamps to accurately trace the steps tried by any individual attempting penetration.

Another consideration is that the SNL Internetwork will be made up of multiple sites that are in different time zones (e.g., Albuquerque and Livermore). These time changes can cause confusion to both software and users and means should be provided to automatically compensate for such differences.

This document does not recommend any particular implementation for the time service in the new networks. It does recommend that the network architecture implementation include a time service that can provide a centralized time. Additionally, or alternatively, the time service could poll the time from the network nodes and maintain the differentials between the internal network clock and the node times obtained.

5.4 Network File Transport

The network file transport service should provide the general capabilities to:

1. Move files from one SSN node to another.
2. Move files between SSN nodes and remote network nodes.
3. Support text and binary files.
4. Provide notification of correct delivery of all data to their final destination.
5. Provide for data tagging to identify the origin of all files.
6. Provide information about the physical characteristics of the transported files (e.g., security levels).

The network file transport service should provide, subject to access control restrictions:

1. Directory services

2. File push capability

3. File pull capability

If feasible and within security restrictions, the network file transport service also should provide:

4. "Third party" file transfer (i.e., user on node A requesting a file transfer between nodes B and C.)

5. File manipulation by user programs (i.e., programmed open for read/write, record access, etc.)

The network file transport service should provide a) high speed (50-100 Mb/s), low transfer-completion-time service between tightly linked computing systems, and b) medium (1-10 Mb/s) transfer-completion-time service between less tightly linked computing resources. A user should be able to logout after initiating this service and expect assured delivery with good error recovery and completion notification indicating success or failure.

Industry standard protocols should be used for compatibility. Only when extenuating requirements such as very high performance are unattainable with standard protocols should proprietary protocols be used.

5.5 Gateways/Bridges

One function of gateways and bridges is to provide connectivity between networks that are within the same security partition (e.g., the SSN and its distributed computing networks, or the PON and its distributed computing networks). In addition, they may also provide connectivity between different distributed computing networks and can perform routing within a single distributed network. Access across information partitions should be through separate "security filters" and not through the gateways.

5.5.1 Gateway Functions

Gateways perform the following useful functions in the interconnection of networks:

1. Operational Isolation. This should permit the SSN/PON and each of the distributed networks to be managed independently. Node addresses within each of the networks can be assigned and nodes can be added, deleted, and moved without impacting the other networks. Also, the impact of errors within a given network can be minimized with respect to the other networks.

2. Security. The routing of data between networks should be permitted only if security policies are satisfied. Gateways should perform node validation and, for secure networks, should be able to associate a range of allowable security levels, and possibly compartments, with each node.

   For a secure network, if a network protocol is capable of labeling connections and/or packets with a security classification, then those labels should be checked against the allowable ranges for the source and destination nodes, and the connection should be disallowed if the label falls outside those ranges. If a network protocol is not capable of supporting individual security labels, then a single security classification should be presumed for that subnetwork, and security checking should be based upon that single level.

3. Protocol Conversion. Currently the VAX Gateway connects a rather extensive DECnet network of VAX/VMS systems using DECnet protocol to the existing CCN, which uses the MASSNET protocol. Future gateways (or enhancements to the existing VAX Gateway) must support additional protocols, which will most likely include NETEX and TCP/IP. In addition, it is anticipated that a potentially large network of UNIX systems and workstations running TCP/IP will need to be connected to the SSN/PON. Eventually, other protocols such as OSI and FDDI will need to be supported by gateways.

   Protocol conversion may be done at several layers of the protocol, depending upon the service being provided. If the same protocol is supported in both a distributed network and the SSN/PON, as would happen if TCP/IP were supported in a workstation network and in the new supercomputer, then the gateway would function primarily as a packet router, although it would still perform the security functions. For cases where protocol compatibility does not exist between the SSN/PON and a distributed subnetwork, the protocol translation should be done at the transport and/or application layers.

4. Buffering/Transmission Rate Matching. The connected external distributed networks will have a wide range of transmission rates that the gateways will have to accommodate. Individual line speeds coming into the gateways will range from 1 to 100 Mb/s with the actual average application throughput being much less than the 100 Mb/s. In order to maintain the performance of the higher speed networks, the individual gateways may require large memory capacities to buffer the data transfers between the networks over this wide performance range. Additional gateways with multiple high bandwidth channels capable of supporting bandwidths of greater than 50 to 100 Mb/s may be required to support the aggregate high end performance. In order to achieve the high bandwidth performance through the gateways, it is not recommended that files be staged to gateway disks. The elimination of staging implies that the network file store and the Supercomputer-90 system both will need to support a large number of open connections and concurrent file transfers. In order to prevent gateway overload, throttles can be built into the gateways to limit the concurrent number and size of file transfers to or through any given node.

5. Checkpoint/Restart and Error Recovery. A gateway is vulnerable to errors on either of the networks it bridges. This adversely affects the mean time between failure. Therefore, a gateway needs significant user-transparent restart capabilities and robust error recovery.

5.5.2 Gateway Services

Gateways should provide few, if any, services of their own. Rather, they provide access from nodes in one network to services in another network. The gateways should ideally provide access to all of the services in the SSN/PON from any node in the supported distributed computing networks. However, it may not be practical to develop and/or support the software to provide the complete matrix of services. The following list of proposed gateway services is given in priority order, with the expectation that all services will eventually be provided. Note that the term "distributed system" is used to describe both multi-user distributed timesharing systems and single-user workstations.

1. Access to the user authentication service in the SSN/PON. This may include capabilities for a user on a distributed system to cause his SSN/PON password to be used on that system. The purpose is to minimize the number of unique passwords per user within a given security partition.

2. Network virtual terminal access from user terminals and workstations to the SSN/PON workers (e.g., supercomputers). This is an extremely important capability, as it is likely that the traditional RS232 terminals will be augmented by PC's and workstations that are connected via LANs into various distributed computing networks that are further connected to the SSN/PON.

3. Access to timer services in the SSN/PON, so that distributed computers and workstations can synchronize their system clocks.

4. File transport between SSN/PON workers (e.g., supercomputers) and distributed systems. Both text and binary files should be supported.

5. Access to batch job queuing facilities for SSN/PON worker computers.

6. Access to a network file storage service. The IFS currently provides this service, but any replacement file storage must be supported as well.

7. Task-to-task connections between processes on distributed systems and processes on the supercomputers. This capability could, for example, be used to distribute computations between supercomputers and workstations, with number crunching being performed on the supercomputer and graphics rendering being performed on a workstation with specialized hardware. An example of this is X-Windows, which uses TCP/IP task-to-task communications to transport graphics data to a workstation window.

8. Access from distributed systems to specialized output processing facilities in the Central Computing Facility, such as microfiche, 35mm and 16mm film, high-volume laser printers, etc. It is anticipated that most output will be directed to distributed systems that are located near the users, since many low-cost high-quality laser printers are available for low and medium volume output. Also, many of the departments doing a large volume of computing probably will purchase their own output devices for video recording and color hardcopy.

9. Electronic mail forwarding between networks, in cases where software is available to support the mail protocols and transport protocols. It is expected that the recently adopted ANSI X.400 mail standard will facilitate the support of internetwork mail.

10. Network interfacing between different distributed LANs. This may require simple packet routing if the distributed networks are running the same protocols, or it may require protocol translation if they are not.

5.6 Batch Job Scheduling

The batch job scheduling facility should allow a network user to submit batch run streams to a compute server such as the new supercomputer, monitor the progress of batch run streams, and exercise some control over the run streams which have not completed. Batch job scheduling should be available on any SSN/PON system which offers batch execution facilities.

1. It should be possible to schedule batch jobs on any network-connected resource which can support large-scale batch execution (e.g., CRAY, IBM, possibly VAX). This may require a phased implementation and interfacing of several batch job scheduling systems.

2. Batch job scheduling should be possible between systems where the work consists of multiple batch jobs. For example, it should be possible to submit a batch stream which executes a scientific job on a CRAY, then submits a batch job to copy the output to, and post-process the data on, a workstation.

3. Batch job scheduling should work with multiple, industry-standard communications protocols, either natively, as a standard, site tailoring option, or via a gateway. The option selected for interfacing a communications protocol should not necessitate source code modifications to the base batch scheduling system.

4. The batch job scheduling system should support the following features for users:

   a. Submit a batch run stream consisting of one or more batch jobs.

   b. Check job status and send the status information back to the terminal.

   c. Cancel jobs waiting for execution or in execution.

   d. Purge job output waiting to be processed or cancel output processing in progress.

   e. Cyclic scheduling. That is, retain a job and automatically resubmit it at various times or intervals.

   f. Precedence/successor relationships. That is, the user should be able to submit a batch run stream which will execute the jobs in it in a specified order with dependencies (i.e., submit job A after job B has run successfully).

   g. Conditional execution. That is, jobs are scheduled only if certain preconditions are satisfied, such as the successful completion of a prior job, or the availability of a network partner, etc.

5.7 Network File Storage Service

The network file storage service is to provide long-term on-line, archival, and backup storage of data. A secondary purpose is to support a temporary data storage overflow condition on worker or service nodes. It is within the realm of possibility that a need for this service may begin to stabilize or decline as cost-effective, high-volume optical storage becomes more readily available to distributed LANs.

1. Transfer Rate. To support SNL, Albuquerque's Supercomputer-90 system, the network file storage service must be able to move one file between the file storage system and the supercomputer at a user-perceived file transfer rate of 50 Mb/s. The minimum transfer rate with other nodes should range between 1 and 10 Mb/s. This significantly lower performance requirement for the non-supercomputer nodes will be dictated by the specific node performance capabilities and/or the presence of Ethernets in the network communication topology. The total (user perceived) throughput of the system is to be 100 Mb/s. (See Section 10.0 for further discussion on network performance.)

2. File Size. A logical file size maximum of 2.1475 gigabytes (2^{31} bytes) is required. One reason this size was chosen was because it will hold a complete memory dump from a supercomputer with 256 megawords (64 bit) of memory. The primary recipient of files of this size will be the network file storage service; therefore, the design and implementation of the new network file store should support both logical and physical files of 2.1475 gigabytes. The degree of the requirement for the other network nodes such as gateways, output, etc. to support this file size is not currently known. Therefore, as an interim measure, in order for other support nodes to accommodate logical files of this size it may be necessary to divide the logical file into smaller physical pieces. (We currently support a logical and physical file size maximum of 150 megabytes). As the requirements are established for the other support nodes to accommodate 2.1475 gigabyte size files, they may need to be upgraded and expanded.

3. Software. The basic functions the software must support are:

   a. Storage of a file
   b. Retrieval of a file or a partial file
   c. Open of a file and random read or write of specific records
   d. Display of file information
   e. Deletion of a file
   f. File management tools (e.g., rename, add subdirectories, etc.)

   Additional desirable features are:

   a. Wildcard capabilities
   b. Maintenance of file attributes, including machine origin
   c. Maintenance of versions
   d. Group authorization
   e. Archival capability near the facility or with the user
   f. Automatic backup/multiple copies
   g. Offsite backup
   h. Program and non-interactive (e.g., batch) interface

4. Network Interfaces. Redundant paths to high-use nodes (e.g., supercomputers) are desirable.

5. Protocols. To provide access from multiple networks, a need for supporting multiple protocols (e.g., NETEX and TCP/IP) is anticipated. The network file storage service should use the same network file transport system(s) as other systems.

6. Security. The network file storage service for the SSN must support files of different classifications and categories (if implemented).

7. Storage Hierarchy. On-line storage should be available for files that have been recently used. Archival storage should be used for inactive files. A user transparent, automated migration process should move bit files between the on-line and archival storage. A backup capability should be implemented to backup on-line storage and user specified bit files.

8. Migration. The file migration mechanism between the SSN and CCN must easily support the migration of files from the IFS to the network file storage service.

5.8 Task-to-Task Communication

Task-to-task communication refers to the capability to support a logical connection for the exchange of binary or text messages of arbitrary length between user applications running on different network nodes. This implies that the user applications must implement some common application level protocol. One of the significant advantages of task-to-task communication is that it may be used to distribute processing between network nodes so as to take advantage of the unique capabilities of each. For example, a user might perform number-crunching on a supercomputer and pass the results in real time to a task running on a graphics workstation which is post-processing and displaying the results.

For secure networks, the message exchange must conform to the security policies established for the network and will, therefore, need to operate under the same basic constraints that are established for other network services such as file transfer and virtual terminal login. Most likely this will require that logical connections be limited to system registered user applications that are running at the same security level and compartment(s) (if implemented) and under the same user ID. End-to-end encryption of the task-to-task logical links is not a firm requirement, especially since approved hardware/software to perform this encryption is not readily available now, but may be desirable if and when it does become commercially available.

Since this service may be a means of introducing Trojan horses/viruses, only "certified, registered tasks" should be allowed to communicate in some environments. The certification and registration of the tasks would be performed by the implementer and be approved as part of the overall security plan.

The features that are desired for task-to-task communications are as follows:

1. A task-to-task subroutine library should be supplied that provides a common application level interface to the task-to-task facility, independent of the network protocol and operating system. This library should implement a simple application protocol to support some of the task-to-task features listed below.

2. Message lengths of any size should be supported, subject to memory limitations. If the application message size exceeds the allowable network message size, the application protocol should segment application messages into multiple network messages and then reassemble them on the receiving side, transparent to the user application.

3. Both text and binary data should be supported. Data type indicators should be passed in the application protocol and should be made available to the user applications. The data type indicators may be used by the applications to determine how (or whether) to translate the data being received from the partner.

4. Effective user data rates of 1 to 10 Mb/s should be possible in both directions, if the user applications are properly designed, i.e., if they do not require excessive handshaking, protocol conversion, or security processing.

5.9 Output Processing

There will continue to be a role for a centralized output service. Output devices are becoming more plentiful in the user work areas, but a network service that provides high volume and special function devices is needed.

The current Output Node provides this service for microfiche, 35mm and 16mm film, video animation, laser plotting, 8x10 color, and high-volume laser printing. The only additional services that have been identified as a requirement are a high-speed two-sided laser printer using cut paper and support for "very high speed graphics".

The support of "very high speed graphics" (Ultra frame buffer type of graphics) should be supported as a separate node due to its unique (800 Mb/s) performance.

An output processing service should be implemented as a Protected Network Resource (i.e., trusted component with no user code and no capability for user access). The goal is that all classifications and network partitions can use it. Alternatively, the service must be partially or completely replicated for disconnected networks which need the output service.

Even assuming that much of the current output node software is reusable for the new service, some enhancements to that software will be needed to accommodate the new network architecture. The output service needs to be reachable, either directly or indirectly, from all networks. Additional protocol support is likely to be needed.

5.10 Network Monitoring and Status Reporting

The operational monitoring of the network will be accomplished by an independent support Management and Operations subnetwork. It will provide the facilities to collect node and network data and information, analyze and display critical information, and may ultimately be able to support network and node diagnostics. Additional information on this subnet is contained in Chapter 6.

5.11 Network Mail and Messaging Facility

Network mail and messaging include the following facilities:

Broadcast	High-impact, immediate information, limited to several lines of data.
News	Information of general interest to network users concerning the use and operation of the network and its facilities, limited to one or two screens of data. News is provided upon user request.
Mail	Peer-to-peer (users or programs) information, limited to the equivalent of several pages on average.

Since mail and messaging facilities on the SSN/PON are not meant to be used as full-function file transport mechanisms, transfer of non-text (i.e., binary) will not be supported.

1. The mail and messaging facility should work with multiple, industry-standard communications protocols, either natively as a standard, site tailoring option, or via a gateway. The option selected for interfacing a communications protocol should not necessitate source code modifications to the base mail/messaging system.

2. Network users should have the option to disable receipt of broadcast, news, and/or mail.

3. All network processors should share a common news/broadcast data base, which is centrally maintained. Due to operational considerations, ordinary network users should not be able to enter broadcast or news items.

4. The mail system should support network, user-to-user communication, including the use of distribution lists. A network application may also be a user for this purpose.

5. Mail (i.e., peer-to-peer information) should be sent end-to-end (from source to destination) and not retained or journaled centrally, even though records of a transaction may be retained centrally.

6. It is desirable, but not mandatory, that the mail system provide the following features:

   a. A supported interface to industry-standard mail facilities, such as All-IN-1, PROFS and VAX MAIL (X.400).

   b. A calendar and "tickler" file for network users. Ideally, these files are associated with the mail system, and not any particular host. The mail system should determine when these files exist for a network user and make them available for display or update.

7. Each host which supports network access should provide network users the ability to perform electronic mail functions, such as sending/reading mail, modifying "tickler" files, etc.

8. It is desirable that a mail system provide an automatic mail forwarding capability so that a user may collect all mail on a single machine in a partition, rather than be forced to check mail on several.

9. The mail and messaging facility should provide connectivity between security partitions, within the rules and constraints for exchanging data between partitions. It should ensure data privacy and, in a secure network, provide for the exchange of both classified and unclassified information.

5.12 Accounting/Billing Service

The accounting/billing function in the network will be limited to collecting and formatting the required accounting data in order for it to be provided to the administrative partition for the preparation of bills. The data may be collected at each individual node or may be centrally collected by the Management and Operations node. The specific implementation will depend on the specific resources available to implement the accounting data collection and formatting function.

6.0 Network Architecture Topology

The architecture proposed for the new network is based on the implementation of three new physical networks that will co-exist with the current network until it is removed. A file migration capability will connect the old and the new environments. There are some general architecture and design philosophies that apply to all of the elements of the new design.

1. The old CCN will remain functional and be connected to the new SSN through a file migration facility. It is not intended that the current implementations of network services be made available in the new environment since the new environment will be UNIX based and has a functional set of capabilities equivalent, but not identical, to the old network services.

2. The new architecture consists of three physically separate information partitions--secure, administrative, and open--each being implemented as a separate physical network or networks. The motivation for the separation is based on two primary considerations: system functionality and security. Connections between the physical networks will be accomplished through "filters" that provide the required control and management over the movement of information/data. Each information partition network can be implemented and tuned to provide optimum functions to its own user community. It is envisioned that the administrative network will be implemented as an IBM SNA network and that the secure and open networks will be implemented as combinations of HYPERchannel and Ethernet networks.

3. Open and secure networks will be implemented with multiple physical subnetworks. Multiple physical subnetworks within an information partition will provide redundancy and increased capacity. As a consequence of the redundancy, we will achieve greater reliability and the ability to match subnetwork bandwidths and protocols to individual node performance requirements. Multiple physical subnetworks will also form the basis of a network migration strategy to accommodate new protocols and networking hardware. And finally, they will provide for the introduction of multiple vendor networking hardware and software products.

4. The backbone secure subnetwork will be implemented with NSC "N"-Series Adapters and a fiber optic physical plant. The "N"-Series Adapters are expected to provide the initial required bandwidths (100 Mb/s) as a supported network product for the spectrum of machines that will be used in the secure network and will maintain compatibility with the current CCN backbone hardware. Additional subnetworks consisting of other 10 to 100 Mb/s networks will be implemented to support unique network functions. The open information partition will be implemented with either "A" series NSC adapters (50 Mb/s) or "N" series (100 Mb/s), depending on the node performance requirements.

5. Within each network, the implementation of services and protocols will be modular to allow for the replication of service nodes and subnetworks to provide increased capacity.

6. All functions specified will be implemented as separate, modular and independent functions in order to allow them to be implemented in a single node or in a number of separate nodes.

7. The backbone network(s), the terminal subnetwork(s), and the security subnetwork(s) will each require a connection to the management/operation subnetwork(s) for network operation.

8. The backbone network(s), the terminal subnetwork(s), and the management/operation subnetwork(s) will each require a connection to the security subnetwork(s) for network security monitoring and operation.

The following figures and discussion describe the new network architecture. Although only one service node may be illustrated in the figures, it is intended that additional "like" services nodes can be implemented as needed. In addition, although we have labeled specific nodes to contain specific services, this is intended to be a generic illustration of how the functions will be delivered. Some of the specific functions can be implemented on a single node while others can be implemented on more than one node. The minimum sets of physical subnetworks have been illustrated. Additional physical subnetworks can be implemented to satisfy special requirements, e.g., high speed graphics and connections between supercomputer workers and storage nodes.

Figure 6.1 shows the interrelationship of the SSN, ADNET, and PON. The direction of the arrows in the figure indicate the direction of possible automated information flow between the networks. The three networks shown are all distinct networks interconnected with "filters" that manage and control the flow of information between the networks. These "filters" may be implemented as either an administrative procedure, a tightly controlled electronic path, or a combination of the two that form a well defined set of rules for the movement of information.

Figure 6.2 illustrates a candidate implementation of the "filter" connections between the three networks (the filters will be developed and implemented as a separate project). All information leaving the secure information partition should be tightly controlled by administrative procedures. The information will probably be moved as either hard copy or removable magnetic media. The information flow between the Open and Administrative partitions might be accomplished by dual-ported disk with appropriate administrative control. The movement between Administrative and Open to Secure could be by file pull from the secure partition, by means of a dual ported disk and encryption. These proposals are being made as a means to provide the required separation of information and the required audit and control of information movement between the different information partitions with different associated risk factors.

Figure 6.3 shows how the CCN will connect to the SSN. They are separate entities with only a distinct two-way file communications path connecting them. The proposal to implement the new network as a separate network with only a migratory communications path between the old and new networks is critical to the development of the new environment with a minimum amount of disruption to the currently running network. Figure 6.3 basically shows the current CCN with the currently installed "A" series NSC adapters and 50 Mb/s coax trunks. This network would remain the primary user network while the new network is being developed and installed. When the new network becomes operational, the File Migration facility would be used to communicate between the two networks. Portions of the current CCN would remain operational for a considerable period of time to satisfy compatibility requirements, such as the need to maintain CFS for use by NWCNET.

Figures 6.4 and 6.5 show the general topology of the SSN. The figures illustrate a minimum configuration for the number of physical subnetworks. The SSN should be composed of three major types of subnetworks: 1) a high bandwidth backbone subnetwork composed of two to four 100 Mb/s fiber (FDDI fiber plant) trunks, 2) a number of dedicated point-to-point 100 to 200 Mb/s subnetworks, 3) three high speed (100 Mb/s) support subnetworks. The number of specific subnetworks, network trunks, and adapters may be replicated to provide increased capacity. It should be noted that, although only the SSN is illustrated, the PON could also be implemented employing the same architecture and topology with the specific number of subnetworks, trunks and adapters being adjusted to provide the capacity and redundancy required.

The backbone shown in Figure 6.4 will be a bus architecture based on the NSC "N" series adapters (see Section 6.1). Each node on the network needs to have a minimum of two adapters connected to provide redundancy, and in certain cases such as for the primary compute and file servers, increased capacity. There should be two to four 100 Mb/s trunks to provide the needed connectivity between the nodes. Each node should be capable of supporting multiple protocols with the specific protocol(s) supported being dependent on the nodes' functionality. The initial protocols to be supported on the backbone will probably be TCP/IP and NETEX. New industry standard protocols will be accommodated as they become available.

The dedicated 100 to 200 Mb/s point-to-point subnetworks are intended to provide high bandwidth, dedicated connections between nodes on the network that have unique performance requirements. Examples of this type of connection would be the connection between a Supercomputer-90 worker and the network file storage service and the connection between the Supercomputer-90 worker and a high speed graphics frame buffer. The protocols utilized on these dedicated networks need to be optimized to support the high bandwidth requirement. The candidate hardware that can be utilized for this purpose are the multiple NSC "N" series adapters with 100 Mb/s fiber connections and the Ultra 800 Mb/s channel to frame buffer connection. It is anticipated that other Gb/s range networking hardware will become available over the next 1 to 2 years.

The three support subnetworks shown in Figure 6.5 are intended to provide implementations of three network functions with optimized hardware, protocols, and software. Two of the subnetworks, the Operations/Management and the Security subnetworks, will provide critical support to the operation of the rest of the subnetworks. All three subnetworks are intended to be separate subnetworks that operate at bandwidths of up to 100 Mb/s, depending on the specific node hardware selected. Like the backbone, these subnetworks should support multiple protocols. The initial protocols will probably be TCP/IP and DECnet, with emphasis on TCP/IP. Other industry standard protocols will be implemented as they become available.

The Operations and Management subnetwork is to 1) provide a path whereby the individual nodes can report their operational status, 2) be used to provide network hardware configuration management data, and 3) be used to initiate network node diagnostics, without interfering with the other operational networks.

The Security subnetwork should be an independent subnetwork that will contain no user traffic. This subnetwork can provide a dedicated path to support both node and user validation and password management, and can provide a path for an independent check on the hardware/software configurations. Because it is an independent subnetwork with a medium bandwidth requirement, perhaps the added overhead of encryption could be accommodated for added security, if necessary.

The terminal subnetwork is intended to provide a path that is optimized to support terminals and other devices that operate with short message or character type traffic. Hardware and protocols optimized for this type of traffic would be used. This subnetwork is needed because short message traffic, when run on a high bandwidth backbone with a protocol optimized for large packets, can severely reduce network efficiency and throughput. It is intended that the routing of traffic onto this network would be by means of a gateway or router that had some knowledge of the speed of the connection (e.g., 9600 baud) or a service like TELNET that would have traffic that would fit the "small message" traffic characteristics.

6.1 Hardware

The hardware for the backbone will be the NSC "DX" technology "N"-series adapters having a 400 Mb/s total backplane bandwidth and providing up to 16 MB of buffer memory. These adapters can provide network interfaces for a 100 Mb/s FDDI fiber plant and 50 Mb/s coax cables, as well as channel interfaces for the variety of machines we might install. The adapter currently provides firmware protocol support in the adapter for NETEX. Firmware protocol support for TCP/IP is still being considered and it is anticipated that support for a full FDDI implementation will be available within 1 year.

6.2 Standard Host Interface

In the development of the future network it is assumed that the operating systems running in a majority of the connected machines will be evolving toward a UNIX environment using the TCP/IP network protocol. The implementation of such a UNIX environment brings with it a host interface that is unique to the UNIX environment. It is therefore assumed that the standard host user interface that will be supported will be based on UNIX and/or POSIX.

The goal of a standard host-user interface is to provide a common command set and a common file structure for as many of the network machines as possible. The benefits to be derived from this approach are portability of applications from machine to machine and improved productivity due to upward compatibility of applications and utilities.

There are disadvantages to consider also. Basic security features currently are not well developed in UNIX, but some improvements are on the horizon. Also, viruses can infect an entire network much more rapidly in a common operating system environment. Networked heterogeneous operating systems have at least some "natural immunity" to being taken over completely by a viral attack. More work must be done in this area to adequately protect interconnected networks using a common operating system from viral attack.

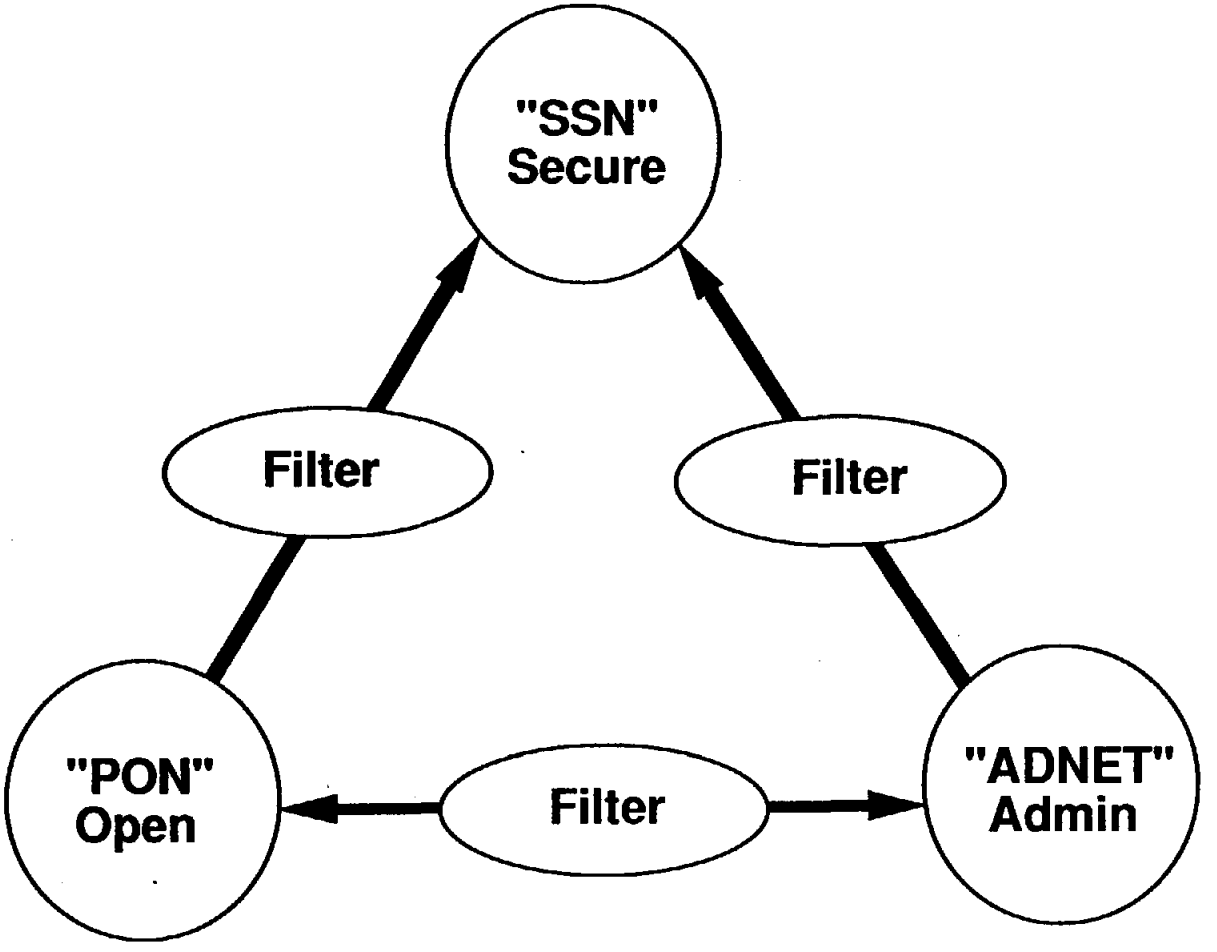

Figure 6.1

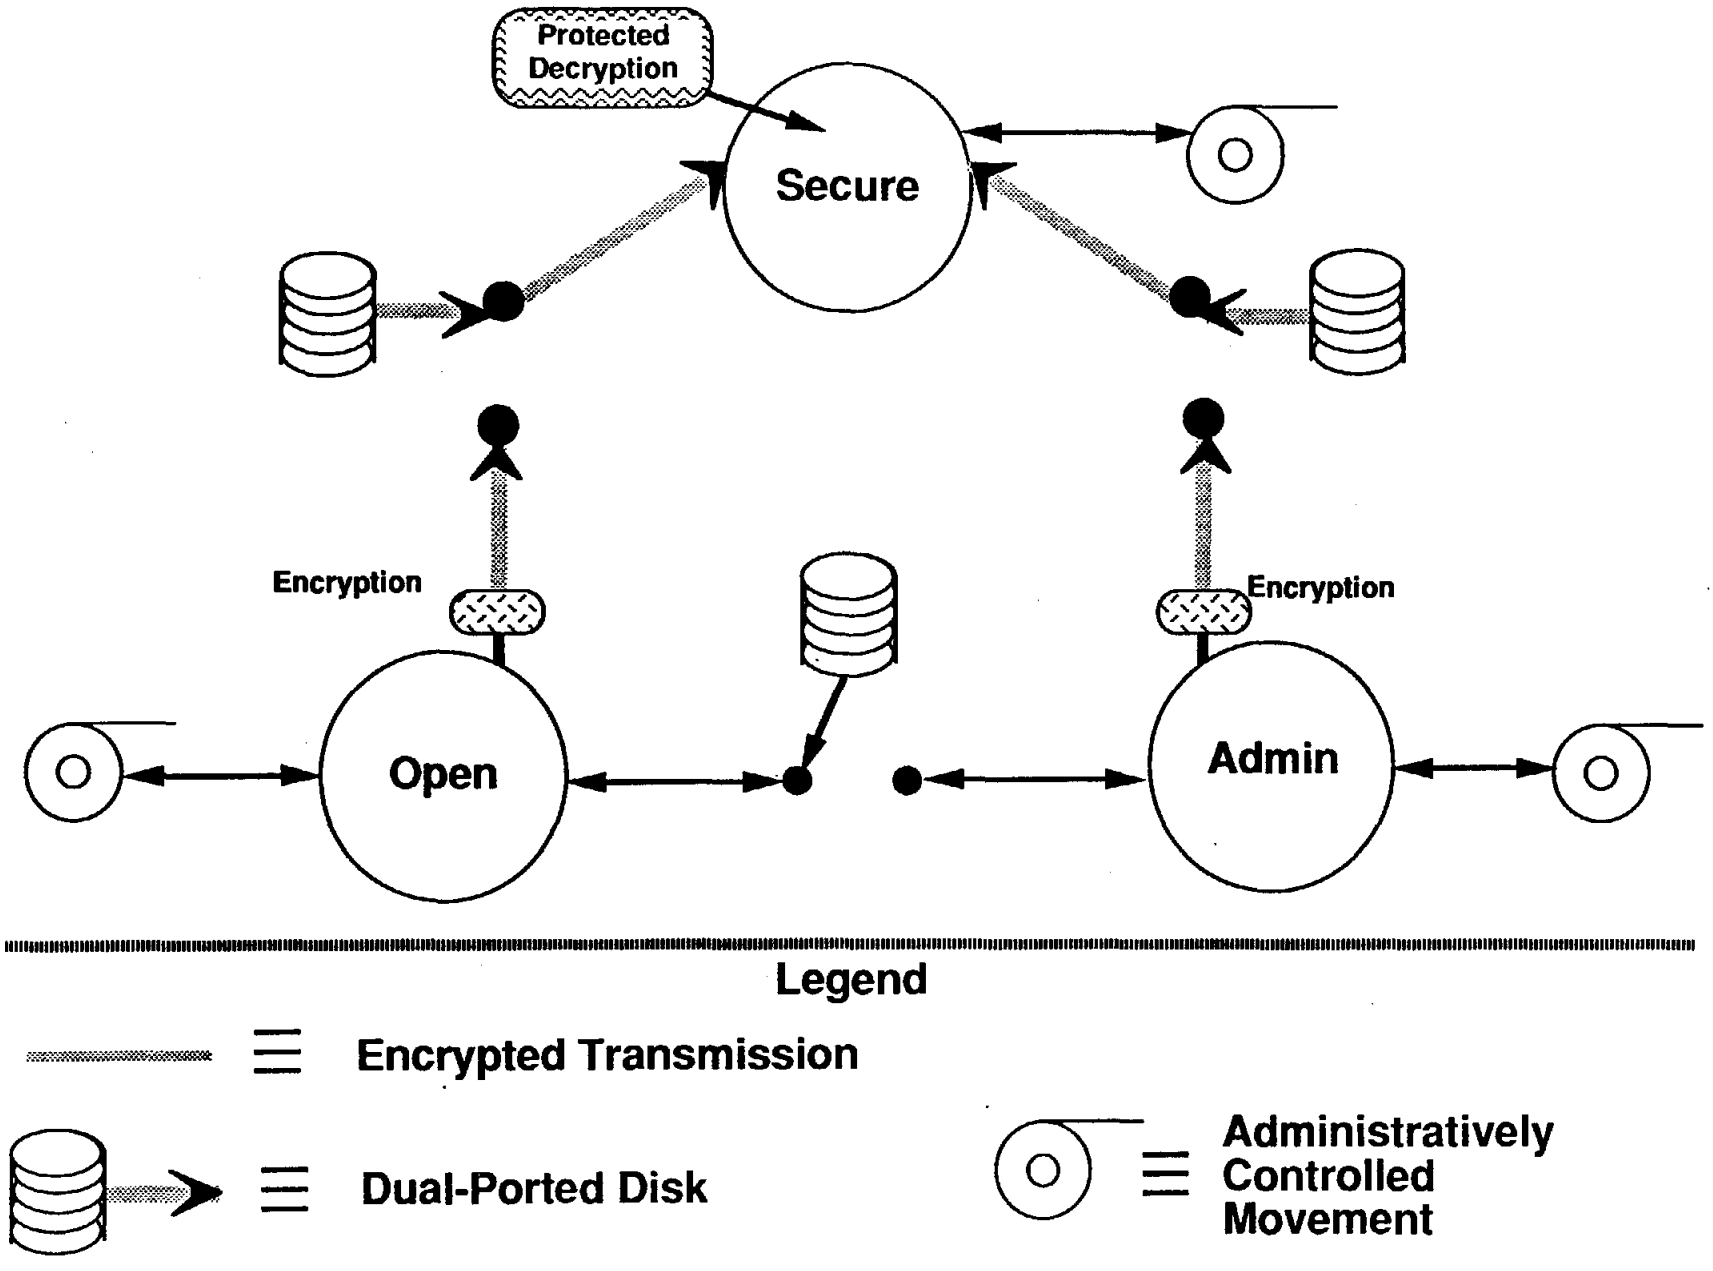
Protected Decryption
Secure
Encryption
Encryption
Open
Admin
Legend
Encrypted Transmission
Dual-Ported Disk
Administratively Controlled Movement

Figure 6.2

Current CCN

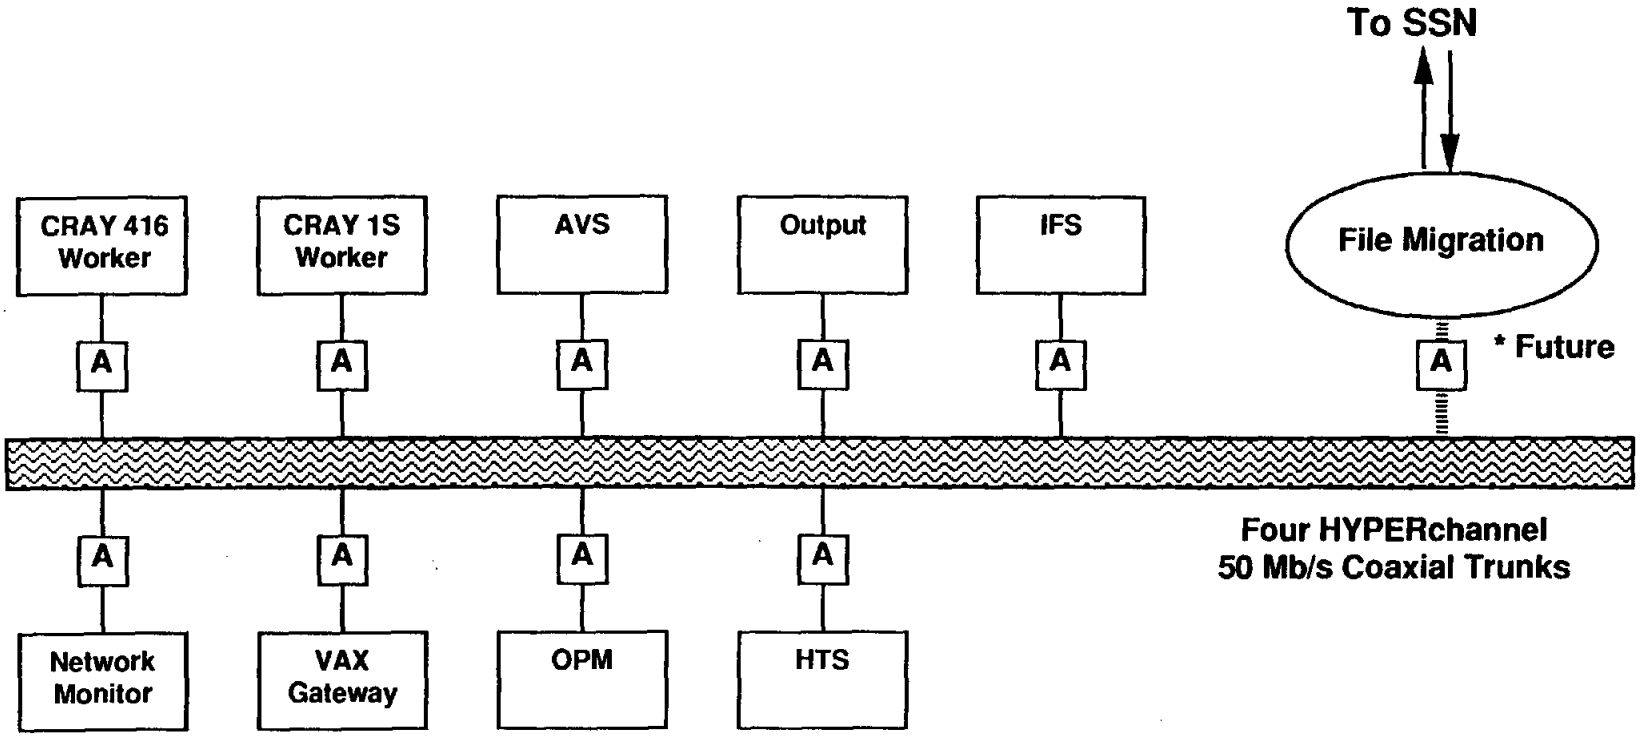

A ⇒ "A" Series NSC Adaptor

Note: Individual functional depictions may consist of multiple physical nodes

Figure 6.3

Backbone
SSN

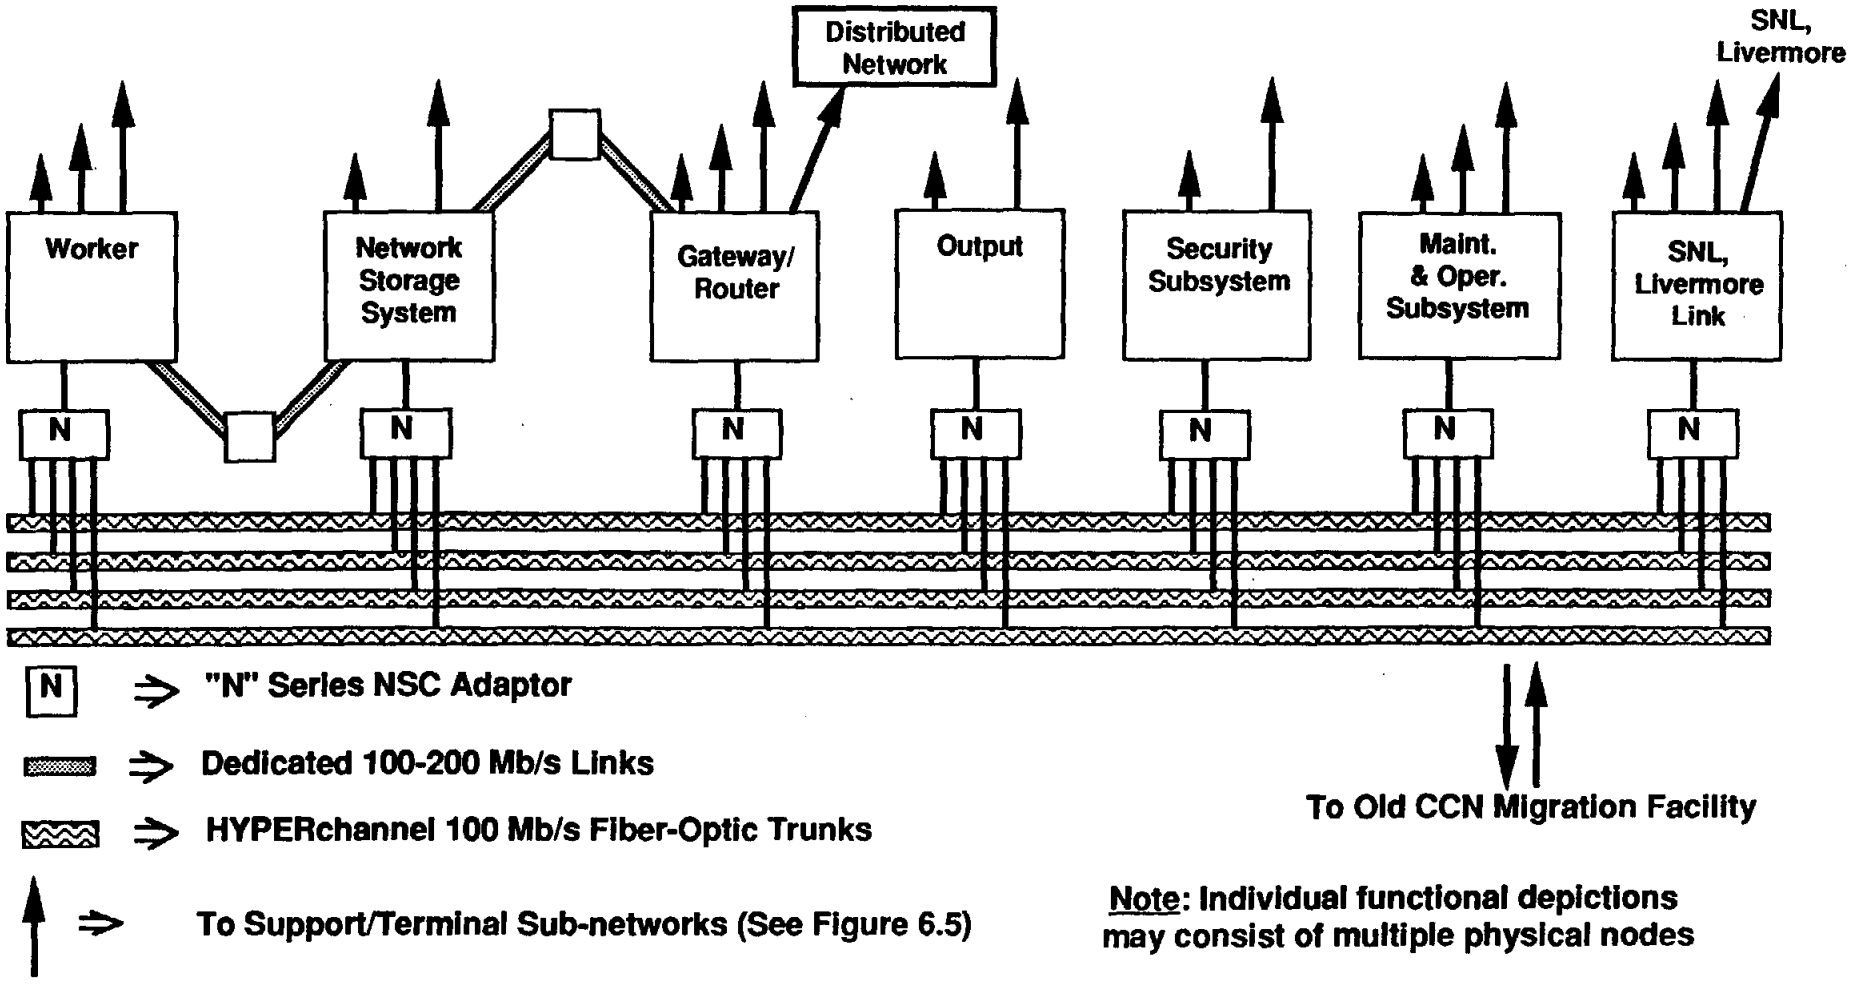

Figure 6.4

Support/Terminal Sub-Network

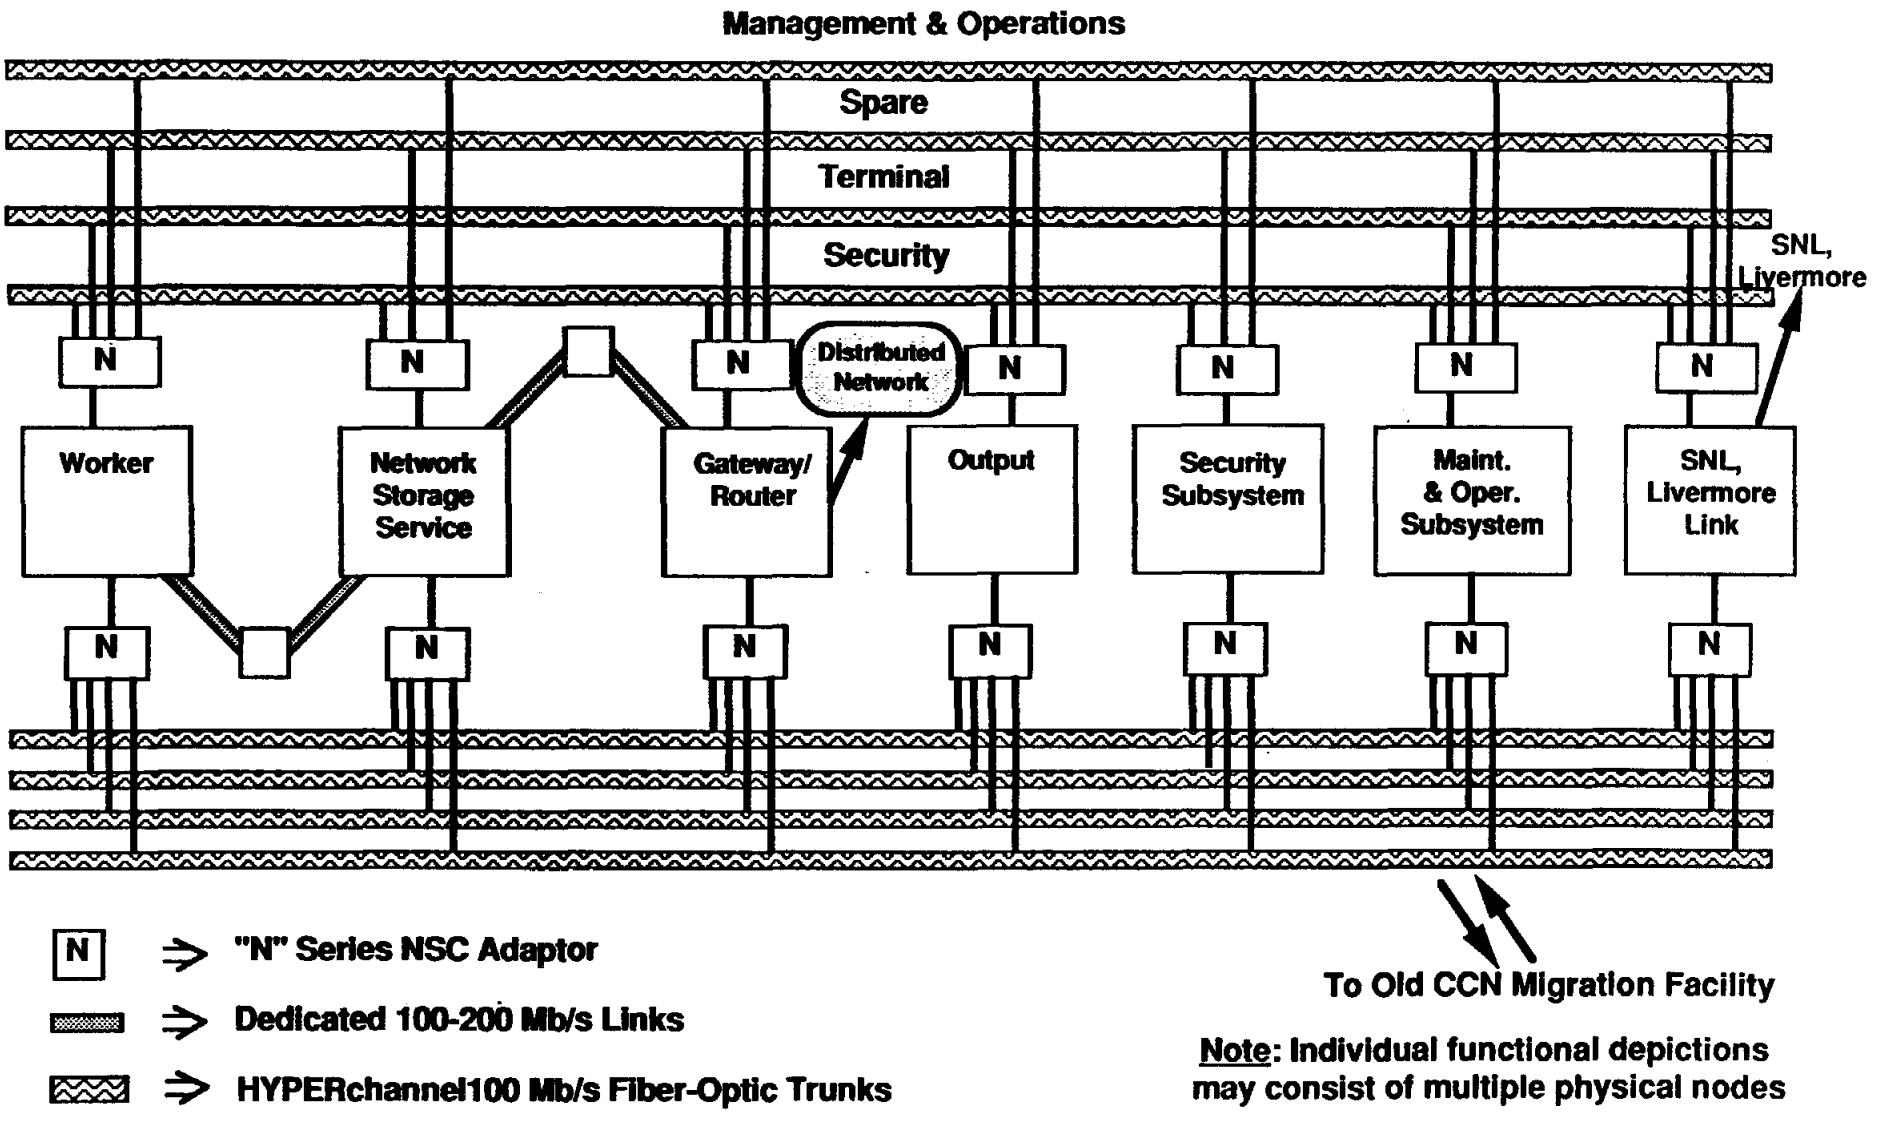

Figure 6.5

7.0 Inter-Network Relationship with: SNL, Livermore; TTR; and other Sites

The SNL, Albuquerque internal networks (SSN and PON) are expected to directly or indirectly connect to similar subnets located at SNL, Livermore and TTR. The capability of data movement to and from similarly controlled networks at other sites will also be provided, subject to additional security controls. The basic intersite networking design philosophy is to allow Sandia personnel at the three major Sandia sites (SNL, Albuquerque; SNL, Livermore; and TTR) to utilize remote computing resources at any of these sites in the same manner that the resource would be accessed locally. The primary goal in the interconnection of Sandia's subnetworks is high transparency of communication at minimum cost. A secondary goal is commonality of usage. These objectives apply to both the SSN and PON.

Transparency of network communication refers to the degree to which a user can distinguish his use of a computing resource from a location local to the computing resource and a location remote from the computing resource. Some of the network differences to be minimized in order to maximize the transparency of the network are: 1) network access procedures, 2) network user throughput, and 3) response time. However, to achieve a very high degree of network transparency may be prohibitively expensive or impossible due to the physical delays, buffering, error control, and flow control problems associated with long distances.

Transparency of network communication also implies that no additional user-visible security controls be imposed between Sandia sites that are not imposed within the Sandia sites. This will not be true of mechanisms to move data between the SNL inter-networks and other similar security partition networks at non-Sandia sites. Additional security controls should be implemented to: 1) restrict access from remote non-Sandia sites to specifically authorized computer resources, 2) disallow the ability to remotely execute user specified instruction strings (no RJE or unregistered task-to-task communication), and 3) allow only registered programs to communicate. A registered task is an application program which has been "certified" to perform only a specific function. These restrictions are deemed necessary to reduce the likelihood of trojan horses/viruses from entering the SNL internal networks via non-Sandia networks. (As subnetworks become interconnected, the personnel exposure increases rapidly, if not exponentially).

Commonality of usage refers to the degree to which a user can distinguish between his use of a computing resource at his local site and his use of a similar computing resource at a remote site. Some of the kinds of differences to be minimized in order to improve commonality of usage are: 1) operating system differences, 2) control language differences, 3) compilers, 4) libraries, and 5) file system structures.

It is expected that the SNL Internet will be connected to the DOE NWCNET by means of removable or switchable media such as a dual- ported-disk. This mechanism can provide the complete NWCNET phase 2 functionality of CFS-CFS file push. CIC#1005 Oct. 88 [Wideband Communications Network (WBCN) Worker Service Reference, Los Alamos National Laboratory] provides WBCN definitions. Sandia may maintain other well controlled electronic connections to other sites to support additional functionality that the NWCNET cannot provide.

Portions of the PON may be interconnected with MILNET, Technet, and the Open LANL XNET. Additional security precautions must be developed for these connections also.

8.0 Protocol Support for the SSN/PON

The following philosophies and design guidelines should guide the implementation of protocol support under the new network architecture. There are four basic ideas to be conveyed: first, the network should implement industry standard or industry supported protocols for the vast majority of the network; second, unique and special purpose protocols should be limited to special applications, such as unique performance requirements or functions; third, the network implementations of the protocols must be capable of simultaneously supporting multiple protocols; and fourth, the implementation of the protocol support should be modular, to ease the migration to new protocols.

1. The anticipated suite of industry-standard and vendor-supported protocols which initially will be used for the SSN/PON network architecture are: TCP/IP, DECnet, and NETEX. Emerging protocols, such as ISO/OSI and protocols associated with FDDI will be supported as they become available from industry, and as the demand arises. It should be noted that this initial protocol suite includes both public-domain and vendor-proprietary, but pseudo-standard, protocols. The ADNET is being designed and implemented as an SNA network, possibly using NETEX to communicate with the UNISYS systems.

2. Although there is a strong preference for the use of industry-standard protocols, it is acknowledged that there may be occasions when the requirement for a non-standard protocol will need to be satisfied. The two conditions under which this is most likely to occur are: 1) when the higher bandwidth of a new hardware implementation is required and a standard protocol is not available and 2) when a required set of features is not available in a standard protocol. These implementations should be undertaken with a great deal of caution since technical implementation support will be limited and continued operational support will have to be provided locally with existing resources. It is envisioned that these implementations should be accomplished ONLY to satisfy a specific and limited requirement. The remainder of the network protocol support should continue to be satisfied by the industry-standard/industry-supported protocols.

3. Vendors should be utilized to provide protocol support. Currently, SNL, Albuquerque provides a significant amount of the support required to evolve and maintain its network protocols. Experience has shown this mode of operation to be expensive, inefficient and inhibitive to the growth of the network and the implementation of new applications. Some of the available network protocols, such as TCP/IP, are in the public domain. However, there is ample vendor support available for industry-standard protocols, whether proprietary or public-domain.

4. Protocols which are proprietary and vendor-supported, but not industry-standard, may be considered if they meet support and implementation requirements. It is likely that some network application requirements can only be satisfied by a proprietary, vendor protocol.

   To be considered as a viable candidate for a network protocol under the new network architecture, the proprietary, vendor protocol should satisfy the following criteria:

a. The protocol shall offer some feature or functionality which is unavailable in industry-standard protocols, but is significant to achieving the objectives of the new network architecture.

b. The vendor shall have a well proven record of successfully implementing the protocol on a large spectrum of hardware and operating environments. Further, the vendor must also have a proven record for successfully extending the protocol to new classes of hardware or operating environments in a timely manner.

c. The vendor shall have a proven record for providing effective, ongoing support (hardware and software) for the candidate protocol.

d. The vendor shall either have provided successfully, or made a strong statement of intent for, an effective interface for the candidate protocol to one or more industry-standard protocols.

An example of a situation that may meet the criteria for a vendor supported, non-industry-standard protocol is the use of NSC's NETEX for HYPERchannel node-to-node transport for functions that do not directly support the user, such as moving a file from the worker node to the file storage node on behalf of the user or system. Although, the industry standard TCP/IP will work over the "A"-Series HYPERchannel, our analysis currently indicates that current implementations might limit performance to values less then or equal to 3 Mb/s on an average non-dedicated network, due partially to packet sizes and protocol handshaking. NETEX, a protocol developed specifically for the NSC hardware, currently will perform at values of 2 to 6 Mb/s on current hardware. Building the new network, utilizing NSC hardware, or for that matter any other specific hardware, with a requirement for an effective minimum bandwidth of 50 to 100 Mb/s will require a protocol that is optimized for that specific hardware to achieve the required 50 to 100 Mb/s performance.

5. When selecting a vendor-supported "special purpose" protocol, a strong emphasis should be given to vendors who currently provide, or have expressed a willingness to provide, security features and interfaces that can support DOE security policies and procedures. Many vendors claim to provide security features in their protocols, but their security approach is rarely compatible with the security required by DOE for networks and network applications. Local implementations of DOE security requirements have been costly, difficult to evolve, difficult to transport to other applications, and very difficult to support.

6. In selecting "special purpose" protocols, a strong technical consideration should be given to selecting protocols which are highly application independent. As networks evolve, there will be continuing requirements to implement applications under a variety of environments, some of which will utilize different protocols. Therefore, a strong preference should be given to the use of protocols which are not heavily embedded in the applications. A standard application-level interface that is straight-forward, simple, consistent, and well-defined would be very beneficial.

7. Protocol translation should be implemented either by gateways/bridges or by vendor-supplied translation facilities. The future network will consist of many subnetworks, each of which may initially use a different protocol. Further, there may be some network applications which must interface with more than one protocol. Hence, there will be protocol translation requirements within the network.

   For inter-network translation, gateways or routers should be utilized. Implementations should be vendor-supplied and "turnkey", as opposed to locally developed. Experience has shown that the development and ongoing support of gateways and bridges is costly and difficult; therefore, more proven industry-supported or industry standard solutions should be implemented.

   Protocol translation at the application level is also difficult and costly, and there is high rate of failure in implementations of this type. Therefore, the use of this type of protocol translation should be minimized.

8. The interfaces for the implementation of the network protocols should be flexible enough to support multiple protocols. As a starting point the three or four protocols referenced earlier would be a minimum set that will require support. It is envisioned that not all nodes will need to support all protocols, but only those that are appropriate to the nodes' functions and locations in the network. The protocol implementation must support protocols that are implemented in the host CPU as well as protocols that are implemented external to the CPU in an intelligent network interface.

9. The generic requirement to support multiple protocols in the new network will be utilized as a basis to phase out obsolete protocols. Once a protocol has been recognized to be obsolete for network use, a plan will be developed to phase out its use. In general, the plan will call for a moratorium on any new application development utilizing the protocol and for a gradual reduction in the level of support for the protocol.

10. The initial set of obsolete protocols for which a planned phase out is required is MASSNET and its local derivatives. While some MASSNET applications may survive for a number of years, it is important to begin immediately to plan for the phasing out of MASSNET. The plan must be explicit and more than just an accepted statement of intent or attitude. Otherwise, MASSNET applications may exist for an indefinite period.

11. Some technique or facility must be provided as an interim measure to bridge between existing MASSNET applications and new applications which utilize industry-standard protocols. The facility or technique provided must be one which may be delivered in a timely fashion and developed in a cost-effective manner. As a rule of thumb, the techniques and facilities should not take so long to implement that users will prefer to develop their own facilities or techniques.

Whenever possible, transparency should be an objective in the development of a bridge. Programmatically, it may be infeasible to do this. In fact, it should be realistic to expect that the bridge, in many instances, will consist mainly of documentation and procedures. Further, it is likely also that users will have to log on and log off of a variety of network facilities to bridge between old and new network applications or functions.

The major criterion for the development of a bridge is that, once developed, it is complete. That is, the user should not be required to develop his own facilities in order to complete the total bridging requirement for a network application or function.

9.0 Network Security

Network security refers to the hardware, software, and administrative measures used to provide the control and audit of data movement through a computer network to prevent the unauthorized disclosure or modification of data. The security policies that will be enforced on the network are those specified for protection of sensitive and classified information by DOE and Sandia management. Of course, the difficulties encountered in attempting to enforce these policies will be much greater due to the presence of heterogeneous hardware and software components in the system and the communications media used to interconnect them. The primary security strategy to enforce these policies will involve implementation of multiple, independent barriers to protect information from access by individuals without proper programmatic authorization and proper "Need-to-Know".

This section primarily addresses the security practices for the secure network. However, good business practices dictate a similar, but less stringent, philosophy in other networks. At a minimum, the specific requirements for processing classified data will not exist for those networks. The security practices for the open information partition and the administrative information partition, as a minimum, need to employ "good" commercial security practices for implementing the SNL, Albuquerque/DOE security policies.

In general, the specific implementations of proposals for providing network/computer security require SNL and DOE approval on a case by case basis. In the following sections, we will describe what we believe are technical solutions, either singly or in combination, to providing adequate network security that would have a reasonable chance of being approved by both SNL and DOE security.

9.1 Network Security Policies

The security policies that are typically enforced on a secure timesharing system are as follows:

1. Users must be validated (authorized and authenticated) before being allowed to access a system.
2. Processes may not access data for which they are not authorized.
3. Processes may not read data that is at a higher classification than the process.
4. Processes may not write data that is at a lower (or higher) classification than the process.
5. Access to classified or sensitive data must be audited, including failed access attempts.
6. Network configurations must be audited.

If we extend the above principles to secure networks of computers, then using the same numbering scheme as the list above, we can derive the following network security requirements:

1. Users and nodes must be validated before being allowed to access the network.

   Nodes must be properly authenticated with the network, so that it is not possible for one node to masquerade as another. This is necessary because, in many cases, user authorization to a given network resource may be based in part upon the identity of the node on which the user is running. Also, the credibility of the network audit trail depends upon being certain of the identity of each node. Node authentication is poor or nonexistent in many network protocols.

   User authentication is necessary because a user's authorization to network resources may be based upon the user's identity. If the authentication performed by a given node is sufficiently trustworthy, then the user might be allowed to access other nodes on the basis of that authentication. However, the most desirable means of performing user authentication is by means of a central network authenticator.

2. Users may not access network resources for which they are not authorized.

   This authorization may be granted by the system manager of each node, based upon the user's prior network authentication. If there is no means for determining that the user has already been properly authenticated, then an additional authentication should be performed by that node.

3. Users may not connect to network processes at a higher classification then what they are currently running.

   Existing network protocols do not currently support any viable mechanism within their structure to limit data communication to be unidirectional; it is, therefore, not possible to enforce a <u>logical</u> link to unidirectional data flow through protocol mechanisms. In order to prevent data from flowing from higher to a lower classification, it is necessary to restrict network logical links to a single classification, with both intercommunicating processes running at that one level of classification.

   If a network node is not capable of labeling classified information and enforcing separation between various classifications, then it should be treated as a system high classification system. Other network nodes should be capable of determining the highest classification and should restrict network logical links with that node to that one classification.

4. Users may not write data that is at a lower (or higher) classification then the process.

   The discussion is the same as 3 above, except for the issue of writing up, which is a factor in preserving data integrity.

5. Network logical links should be audited, especially those created by classified processes. This includes failed attempts to establish network logical links.

6. Network hardware and software configuration management must be audited at "frequent" intervals. The methods and frequency should be determined as a integral part of a total network security design.

9.2 Protocol Support for Security

One requirement that is necessary in order to implement all of the above network security policies, and make them enforceable with a minimum of additional security software, is that the network protocols must be capable of labeling packets with security classifications. This labeling should ideally be contained in each network packet, but it may suffice in some cases to exchange security classifications only at the time a logical link is created. In the latter case, the logical link number contained in each packet would imply the security classification that was established at logical link creation.

Unfortunately, most existing network protocols do not have a capability for labeling packets with security classifications. Two protocols that do have this capability are TCP/IP and the version of DECnet that is released with VMS/SES. The secure DECnet protocol passes complete security classifications when a logical link is created and does not attempt to label each packet. TCP/IP passes security classification labels in each packet at the IP routing layer, but it does not provide any mechanism for the host operating system at the TCP layer to pass the security classification to IP. It will be necessary for either the vendors or Sandia to enhance basic network protocols to provide for the labeling of network packets with security data classifications or to enhance the networking or operating system software to reject logical links if the data classifications of the two cooperating processes do not match.

9.3 Encryption as a Tool for Network Security

Encryption has the potential of being an additional tool for achieving greater network security in the areas of (1) user authentication, (2) node authentication, and (3) packet privacy and integrity. There appear to be several schemes for both user and node authentication, based on public/private key encryption algorithms, that will require further investigation.

End-to-end encryption of network messages at the transport and/or application layers can protect network traffic from being read and modified, even by privileged users at intermediate routing nodes. It could be used to provide need-to-know protection for sensitive and/or classified data as it traverses the network. However, once the data is delivered to a given network node/user, it is up to the local operating system to protect the data from other users on that node (possibly by file storage encryption). The main problem with end-to-end encryption is that it is currently not available commercially. When end-to-end encryption hardware and software do become available commercially and they are shown to be cost effective, we should be able to incorporate that capability into the SNL Internetwork. Until then, we probably will have to trust the routing nodes and restrict privileged access to them as much as possible.

9.4 Separate Physical Subnetworks

One technique that has both positive and negative effects on network security is the separation of the SNL Internetwork into multiple physical networks. This will happen naturally as many organizations build their own LANs and connect them, subject to security constraints, to the SNL Internetwork via Gateways. These LANs need to be managed independently for both administrative and technical reasons, and may each use different communications media and protocols. However, Directorate 2600 (primarily Division 2612) should provide security guidance and administrative coordination of these subnetworks to ensure that they are operated and interconnected in a secure fashion. Gateways operated by 2600's personnel will provide the primary means for interconnecting various subnetworks, including the SSN, into a secure SNL Internetwork.

9.5 Security Filters

Security filters can provide a control and auditing service for transactions which span certain regions (e.g., sites) of the interconnected networks.

Security filters can be installed to restrict or control access to network services and flow of information between regions of the network which have different security characteristics.

The control and auditing service may authorize certain kinds of transactions on the basis of tokens of user authentication and authorization policies established for the operation of the network. In order to prevent subversion of the control and auditing performed by the security filter, the filter mechanism should operate in an environment which executes no user-specified instruction sequences. As part of the security filter design, a covert channel analysis should be performed to minimize the possibility of leakage through such channels.

The security filters should restrict only that functionality which is intended to be restricted for well documented security purposes. Ideally, the permitted services should be transparent to the user. Permitted services should be accessed the same way and provide the same functionality for a user on either side of the security filter.

To minimize the requirements for security reviews of general purpose computers used as security filters, the security filters should be implemented in software/firmware/hardware of dedicated equipment with all unnecessary functionality removed. Alternatively, the filters could be NSA-certified components which have already passed security review.

If developed internally, the filters should be very carefully designed, coded and documented. However, it would be highly desirable to buy a vendor-supported security system that has already been certified at the B3 or A1 level by the National Computer Security Center; several are known to be under development. In that case, the various network security partitions could be connected to corresponding secure, administrative, and open partitions through the A1 operating system, and the system itself could be trusted to enforce the security policy for data movement between partitions.

9.6 Security Partitions

As previously indicated, secure, administrative, and open computing are expected to be partitioned into totally separate networks, both physically and logically. Connections between these networks, if allowed, will be via trusted electronic or procedural security filters, or NSA-certified B3 to A1 operating systems. The hardware and software of these security filters should be subjected to rigorous testing and configuration control.

9.7 Network Security Design

With the development and implementation of a UNIX environment, the requirement for a broad systems level view of the network and all of its connections and elements becomes critical from a security perspective. Any network element and/or connection anywhere in the network can affect the overall security of the entire network. As a first step in addressing this issue, the development of a network security design that addresses the security policies, functional and performance requirements, and the proposed technical solutions is essential. This design should be a complementary effort to the technical functional and implementation designs and the network security plan and test for accreditation.

10.0 Network Performance

10.1 Data Transfer Rates within the SSN/PON

The backbone high speed service is to provide channel-to-channel communications within the SSN/PON. The high speed service is expected to operate at a minimum effective rate of 50 Mb/s per physical connection. This 50 Mb/s rate is based on the 100 Mb/s per trunk raw bandwidth of the standard new network trunks, reduced by the anticipated 50% efficiency of the protocols. The actual user throughput, however, is a complex function of many parameters such as bandwidth, processor power, disk transfer rate, and protocols.

For the dedicated point-to-point networks within the SSN/PON, the minimum effective rate (excluding protocol overhead) needs to be 100 Mb/s. The requirement of 100 Mb/s is based on the following scenario: Assume a supercomputer with a central memory that contains 256×10^6 words, each of 64 bit length, or approximately 1.6×10^{10} bits. If you have a network that has an effective bandwidth of 3 Mb/s, it will take approximately 1.5 hours to move one complete memory dump. If you have a network with a 100 Mb/s effective bandwidth, it will require about 2.5 minutes to move that same 1.6×10^{10} bits. If you assume a 50% efficiency for the network due to the protocols, a raw bandwidth of 200 Mb/s is required to achieve an effective 100 Mb/s.

Some special machine-to-machine connections could be considered that would operate at application-to-application rates up to 100 Mb/s by utilizing one of the emerging 800 Mb/s to 1 Gb/s versions of networking hardware. This would be dependent on identifying and installing a machine, other than another Supercomputer-90 system, that can accept this bandwidth. In general, the maximum rate achievable for this service will be strongly dependent on the maximum bandwidth of the device/channel that the Supercomputer-90 system is communicating with. Another technique to achieve an effective rate greater than 50 Mb/s is to utilize parallel 50 Mb/s connections between the Supercomputer-90 system and the other machines.

10.2 Data Rate Requirements for Supercomputer-to-Workstation Communications

The communications mechanism between the supercomputer and most remote sites is expected to operate at an effective rate of 1 Mb/s to 3 Mb/s (application-to-application) to be able to download graphic images and database snapshots of large codes to workstations. The range of 1 to 3 Mb/s is based on the projected utilization of Ethernets and the PBX currently being installed. Also, most Ethernet host interfaces have bandwidths of only 3-4 Mb/s. (See Section 6.0 for further discussion of the network hardware considerations.)

The following calculations were used to estimate data rate needs for communications between a supercomputer and a graphics workstation. Of course, the amount of data to be transferred for any graphics image depends upon the resolution of the graphics device, the number of colors, and the complexity of the image. Image complexity is significant if run-length encoding is used for a raster image or if a polygon/wire-frame

description of an object is transmitted to a graphics workstation for rendering via specialized hardware. Because there are so many parameters involved in the transmission of graphics images, we shall present the calculations for a few common cases that are representative of near-term SNL, Albuquerque user requirements.

1. Run-length encoded raster image.

   a. Tektronix terminal.

      This is the type of graphics image currently being produced by personnel in Department 1530 using the recently developed Tektronix DMA support in DNTS. They are using Tektronix terminals with 256 colors and a resolution of 1280 X 1024. This calculates to

      8 bits/pixel X 1280 X 1024 = 10,485,760 => 10 Mb/image

      Run-length encoding appears to reduce the amount of data by about an order of magnitude for many of 1530's plots, so that they can paint a screen with 1 Mb of data at a data rate of 200Kb/s in about 5-6 seconds. This screen painting rate appears to be acceptable to the users.

   b. Raster Tech terminal.

      Organization 2600 personnel have been producing movies via a Raster Tech DMA terminal for the past couple of years, and 1200 has recently set up a similar system. The Raster Tech supports 16 million colors (24 bits/pixel) with a resolution of 512 X 512. This calculates to

      24 bits/pixel X 512 X 512 = 6,291,456 => 6 Mb/image

      Thus, the Raster Tech requires approximately the same amount of data per image as the Tektronix. Again, run-length encoding is commonly used, resulting in a screen painting time of 5-6 seconds for a typical plot.

2. Polygon Description of Solid Object

   This approach assumes that the user is sending a polygon description of an object from a supercomputer to a graphics workstation that has specialized hardware capable of rendering the object, and perhaps transforming or rotating the object without further interaction with the supercomputer. This is the mode of usage that may be used commonly in 1-3 years. The X-Windowing System with PHIGS extensions already provides the software for a user application to do this across a network. Organization 1520 personnel have estimated the number of polygons necessary to describe a typical object as being 20K to 50K. Since they are not very familiar with details of the X-Window protocol, they have estimated the amount of data necessary to describe a typical polygon at 40 bytes. This allows us to calculate the amount of data as

20,000 X (40 X 8) = 6.4 Mb/object

50,000 X (40 X 8) = 16 Mb/object

Burst data rates of 2-3 Mb/s, over an Ethernet in memory-to-memory transfers, would enable transfer of the polygon descriptions in a few seconds. These rates are achievable by paying particular attention to protocol tuning and network system configurations.

Organization 1530 personnel have estimated the number of polygon in a typical plot for their types of problems as ranging from 10K to 500K. If we calculate the amount of data in a polygon description of an object at the top end of that range, we get

500,000 X (40 X 8) = 160 Mb/object

At data rates of 2-3 Mb/s, it would take 53-80 seconds (a minute or so) to transfer the polygon description to a workstation, and this probably would be unacceptable to the user. There are two possible solutions to this problem. Either the data rate to the workstation could be increased by an order of magnitude (to 20-30 Mb/s) using technology other than Ethernet, or complicated compression techniques could be employed which might be able to decrease the amount of data by an order of magnitude, so that it would be practical to use lower bandwidth (2-3 Mb/s) channels.

3. Vector Description of a Solid Object

Organization 1522 personnel have estimated that a typical vector description of a 3-D object would require 50K vectors, with 32 bytes needed to describe each vector. This calculates to

50,000 X (32 X 8) = 12.8 Mb/object

Again burst rates of 2-3 Mb/s would allow us to transmit the description in just a few seconds.

11.0 Network Operations/Management & Maintenance

This service is required to determine the operational status of all node interfaces and hosts directly connected on the network. This function may be be provided through the use of a passive (i.e. listen-only) network service to collect, analyze, and report network data. The information obtained from each node should, as a minimum, be sufficient to determine if the network interface and the connected host is operating correctly and to determine the specific hardware and software present on the network. The specific network data to be collected will be dependent on the unique characteristics of each network node, i.e., it will depend on the specific network interface hardware, the specific node hardware, the operating system or micro code and the applications software.

Some of the recommended functionality is specified below:

1. The capability to invoke a limited set of network and host diagnostics to aid in isolating the possible cause of any malfunction is needed. Diagnostics such as interface and host loopback and the initiation of canned host network diagnostics should be available. Periodic end-to-end testing of user application functions from subnet to subnet should automatically be performed and test success should be plotted for monthly review.

   A single well-trained entity should be empowered to resolve conflicts in the diagnosis of complex problems of subnet to subnet application communication. The person should dispatch or coordinate corrective action.

2. The ability to gather and process statistics on retransmissions, effective bandwidth, to/from addresses, total amount of data transmitted, etc. is needed. The statistics should be gathered from each network interface and host. Provision should be made for data to be graphically displayed as a summary of all connected hosts and interfaces as well as by individual interface and host. The capability to archive data for later analysis is needed.

3. The statistics gathering should be controllable to eliminate selected network interfaces and hosts from data collection.

4. Analysis programs should be available to reduce the gathered data to provide predicative indications of the status of the network and should be employed to provide network performance trend analysis.

5. The data gathering and diagnostic software should provide data and information to allow comprehensive network, hardware and software configuration management.

12.0 Network Conventions/Standards

In evolving the CCN from a network that is essentially a central hub into one in which the SSN/PON are subnetworks on a larger network consisting of other LAN networks (e.g., SNL, Livermore), there will be a requirement to develop and implement network conventions and standards. The requirement to develop and implement standards arises from two fundamental requirements; one is to be able to diagnose and fix network failures across the entire network, and a second is the need to be able to install new systems on any of the subnetworks and expect them to operate in conjunction with other existing systems with a minimum amount of modification. If the assertion holds true that the entire environment is moving toward UNIX/POSIX, then some of the conventions/standards recommended by the Scientific Workstation Project [SAND88-1179 Scientific Workstation Project Final Report Nov 1988] would seem to be a starting point for the development of "NETWORK WIDE" standards. We embrace the standards that are proposed by the Workstation Project. An excerpt of their standards is given in Appendix A.

A very important convention is the use of a common language by users and systems personnel. SNL, Albuquerque Division 2614 personnel currently maintain a glossary of many of the important terms. A copy of this glossary, with selected additions, is given in Appendix B.

13.0 Migration Considerations

Once the key components of the new SSN are installed (approximately 2QCY90), users will be faced with migrating from old, familiar applications and procedures to new ones. The migration task could be quite complex, depending on the machines used on the CCN to produce the information being migrated to the SSN. In many cases, it may not be possible to provide the user with automated tools to assist in migration, either because it is infeasible to develop the tools in a timely manner or because the automated tools would be unreasonably complex to build. To help compensate for migration difficulties, the use of the new facilities should be encouraged by significantly improving functionality, performance, and/or reliability. Use of the new facilities already is being encouraged by allowing user access to the Prototype Secure Network (PSN) for testing out much of the new network functionality.

It is not reasonable, in this document, to provide a detailed plan for migrating each application or function. This can be accomplished only on a case-by-case basis, and only after the new application or function has been acquired and studied or is well past the internal design phase.

Specific migration plans and implementations should achieve high-level objectives for migration and should follow some general guidelines to allow an orderly transition to new facilities. The high-level goals for migration are:

1. The use of new facilities should be encouraged. New facilities should include the latest hardware and software, be provided with the best support, and receive the most attention for educational efforts.

2. The use of obsolete facilities should be discouraged. A moratorium should be declared on enhancing obsolete facilities, even before the new facilities are available. Obsolete facilities should be minimally supported and these efforts should be slowly scaled down.

3. The complexity and difficulty of migration should be reduced for the users, to the greatest degree possible. Migration can be a daunting task and there is frequently a great deal of inertia against embarking on it. Automated tools, procedures and technical support should be provided to assist in the migration task. However, the assistance provided toward migration must be limited to that which can be accomplished in a timely manner, with obvious savings or benefits.

   For example, suppose an automated procedure for migrating files from an obsolete to a new facility can be developed to save users from using manual methods. It may be worthwhile to provide a migration aid if it can be developed by a single individual in three months, but may not be worthwhile if it were to require three people for a year. The decision to build automated migration aids must be made on a case-by-case basis.

Some suggested general guidelines for the development and implementation of migration aids are:

1. New facilities should be implemented in parallel with obsolete ones, whenever possible.

2. In the case where the new facility must replace the old facility, the new facility should be run in parallel with the old one to test and verify that it can successfully replace it. Unless the new facility proves unworkable, the obsolete facility should be disconnected at a predetermined date.

3. There should be an extended period where new facilities are being integrated into applications or systems which use both new and obsolete facilities. To reduce the complexity and difficulty of the migration, new facilities should be integrated in a systematic fashion.

4. Whenever a new facility is implemented into an integrated setting, it should not be made available for users until tools or procedures can be developed and tested which allow the user to continue to use the integrated product.

5. Implementations of new facilities into an integrated product should be externally transparent to users, whenever possible. When this is not possible or not feasible, users should be provided with adequate procedures, documentation and education to allow them to adapt to using the integrated product in a new way. Changes of this type must be presented to the user with enough lead time to allow them to plan their own migration tasks.

6. The preferred method of implementing new facilities and phasing out obsolete facilities is through the use of physical separation (i.e., physically separate hardware, networks, etc.). Whenever possible, internal flow between these physically separate components should be transparent to users. When this is not possible or not feasible, users should be provided with adequate procedures, documentation and education to allow them to effect the flow between separate components.

7. The preferred method for discontinuing use and support of an obsolete facility is to simply allow use of the obsolete facility to decline. When use declines to a predetermined level, then a scheduled date for removal of that facility should be announced and the facility should be removed on that date.

8. Some planned educational effort should be undertaken, well in advance of implementation, for each new facility, and even for the new network architecture as a whole. Users must be advised of how the new facilities will function, how they will differ from the obsolete facilities, and what the time table for implementation will be. Further, they should be advised, as soon as possible, of specific migration plans and methodologies, and they must be helped to understand the scope of their work in the migration effort.

14.0 Network Reliability and Maintainability

Improved reliability and maintainability of the computing/networking environment is a critical goal of this design effort. To address this goal, we have planned a five-pronged approach: 1) provide an architecture that provides redundant network interfaces and network trunks, 2) analyze the functional designs and implementation designs to insure that we are aware of the critical system elements and single points of failure, 3) design and install the network physical plant so that it is maintainable (i.e., equipped with patch and test panels, spare media, test equipment, etc.), 4) develop and implement the ability to monitor, analyze, and report on the performance of the network physical plant and network nodes, and 5) install fault tolerant systems for the critical nodes either by buying fault tolerant systems or by implementing the system with primary and backup elements. We believe that, if we effectively address each of these areas in the design and implementation phases of the new network, we will realize an improvement over the reliability and maintainability of the network we have today. In addition, the network data collection and analysis effort will allow us to monitor many more aspects of the system performance and to make better system adjustments then we currently do.

The proposed architecture has addressed the reliability by defining a topology that has redundant network interfaces on each node and multiple networks that can provide some degree of backup. In addition, a network physical plant monitoring system is expected to be an integral part of the design. The dedicated Management and Operations function should provide the collection of network interface and network node data and provide analysis and reporting of the network's status very quickly both to identify network problems and to verify that they have been repaired. The redundant network interfaces and network paths, combined with the network reporting and added fault tolerance for critical nodes, is anticipated to contribute positively to the total system reliability.

15.0 Design Document Revision and Control

This design document is expected to be a "living" document and will be subject to revision as conditions and requirements change. Minor corrections and/or clarifications may result in revisions of the document (e.g., Rev. 1.1, Rev. 1.2, etc.) with a limited distribution. Major changes to the design, for example, as a result of a formal design review, will result in the issuance of a new design document (e.g., Rev. 2.0, Rev. 3.0, etc.) with a wide distribution.

Appendix A. Excerpt from Scientific Workstation Document [SAND88-1179]

Note: In this excerpt the references to CCN, the current network, should be interpreted to apply to the SSN/PON networks, currently under development.

Procedures/Standards/Conventions

Need for Procedures/Standards/Conventions

Without some specific procedures, conventions, and standards which are centrally controlled, LANs would evolve randomly. An uncontrolled LAN is undesirable because it creates support and security problems. Growth can and should be regulated/controlled and planned. Future development, flexibility, consistency, and expandability of workstations/LANs depend upon the level of standardization followed by individual workstations/LANs. If one LAN is to ever communicate with another LAN then certain criteria must be met. Failure to adopt and follow procedures could result in massive LAN reconfiguration to allow inter-LAN communication.

LANs (and workstations) should have a common protocol, unique node (machine) names, unique internet addresses, unique user ids, consistent management procedures/tools, etc. Without this uniqueness and consistency, integrity, communication, management, and other problems arise. To alleviate this situation, standards, procedures, and conventions must be established and controlled centrally by groups or organizations within 2600 and followed by the LAN user communities. User input should be encouraged when standards/conventions are being considered for adoption. Thus, support and standardization is necessary not only to meet the users needs, but also to supply flexibility and the consistency that users require.

Recommended Standards/Conventions

Below are details on the standards/conventions recommended by the Scientific Workstation Project.

Network Naming Conventions

Node Names In general, UNIX based machines and those having a TCP/IP interface require a file (/etc/hosts) that contains the internet address, the primary host name, and host aliases for all nodes with which the machine can communicate. The primary host (node) name is used by specific software as the only name for a particular machine. For example UNIX mail uses the primary host name. Other software matches the first occurrence of the host name, either a primary or an alias. The Telnet protocol uses this first match method. If there were to be more than one node with the same node name, integrity problems and/or errors could result. In the case of Telnet, errors are not seen and the result is an attempted access to the wrong machine. The uniqueness of node names is essential to the operation of workstations and their LAN environment. The following discusses requirements, possible solutions-why some solutions are unworkable, and the Scientific Workstation Project's guidelines for naming workstations:

Requirements (if CCN access is required by the workstation)
- o host names must be unique
- o host names are centrally assigned and controlled

Possible options

- o tied to user-NO-users change too frequently
- o tied to organization-NO-organizational changes would affect database
- o tied to bldg. or room location-NO-workstations are too mobile
- o arbitrary name chosen by user-Maybe-but there is usually little consistency used in naming nodes in this manner, also the node name becomes associated more with the user than the node-which will inevitably cause problems when the user gets a new machine and the original node becomes another user's machine
- o standard format name-Yes-e.g. slv#### where s stands for Sandia, l is a site code (a-SNL Albuquerque l-SNL Livermore), v is a vendor id, and #### is an identifier such as a serial number

Guidelines

A standard format name is recommended that could be used Sandia wide. By including site location each site could administer their own node names independently. A vendor code is useful to determine if there are communication or other unusual qualities (good or bad) that are specific to certain vendors. It is desirable to have the serial number be the same as the serial number in the internet address (see below). Matching these two numbers is possible but may be unrealistic because it would: 1) impose too much restriction in node naming, 2) cause gaps in what otherwise might be a serial ordering, and 3) there would be conflicts with older nodes that need an internet address but already have a node name. (e.g. the distributed VAX nodes-SAV###.) Thus, we recommend not attempting to assign the same number for the node name and the internet address. Although, there is an apparent concern for backward capability with DECnet-VMS which limits node names to six characters, this is not a significant problem because: 1) with the next release of DECnet-VMS this limitation is removed and 2) currently ULTRIX machines can provide the mapping between DECnet-VMS and TCP/IP nodes.

Internet address Just as node names must be unique, internet addresses must also be unique. Recommended requirements and guidelines for workstations within Sandia that would have access to the CCN are the following:

Requirements (if CCN access is required by the workstation)

- o address must be unique
- o addresses are centrally assigned and controlled

Guidelines

The internet address is a four byte numeric value which consists of a network address and a local address. The most common scheme uses four (8 bit) fields as shown below. There is a specific group that defines these fields for communication between nodes that are on specific networks such as the DARPA. Since current security requirements imply that any node that is connected to Central Site cannot be connected to an outside (open) node or network, our commendations/guidelines for the use of these fields disregard external networks and are based solely upon defining specific areas and subareas as needed by Sandia to segregate internal LANs. The classification level provides the administrative capability of distinguishing a node's classification level by its address. Below is an example of an internet address and our recommendation for each field.

Internet address format: field 1, field 2, field 3, field 4
example: 125.2.1.25

- o field 1-area
- o field 2-subarea
- o field 3-highest data classification level
- o field 4-serial number unique within area/subarea

User ID and Group ID numbers are present in the password files on both client and server nodes and are used primarily for file protection and access rights. There are several complications that arise when user IDs are not unique. Within a LAN, if two users have the same numeric user ID (regardless of group), those users map to each other. The two users would have equivalent access to all files which reside on a resource available to both users.The implications are that user IDs across LANs must be unique to maintain file protection and integrity. Although use of an E or N number would be desirable and consistent with our present system, currently they are not valid for user IDs, because the range of IDs generally available does not accommodate their size (limited to 16 bits). In addition to user IDs, group IDs are also used for file access and/or restriction of access. These groups can be helpful and/or harmful dependent upon the care taken to establish procedures for their use. Because of common shared resources, such as data storage, user IDs and group IDs must be controlled to insure their uniqueness and appropriate use.

- Requirements (if CCN access is required by the workstation)
  - o numeric user IDs must be unique
  - o group and user IDs are centrally assigned and controlled

Guidelines

- o Assignment of user IDs when the CCN password is issued
- o Assignment of a set of group IDs for each LAN when internet addresses and node names are assigned

Software Standards/Conventions

Software Standards and Conventions are essential in order to maintain control, flexibility, and consistency of the workstation environment. Recommendations and guidelines are described below. The degree to which the software standards and/or conventions are adopted by a LAN user community would be directly proportional to the support that the individual LAN would receive from groups within 2600.

Communications

Protocols-Currently TCP/IP is the standard for UNIX based machines while both DECnet-VMS and TCP/IP would initially be supported for VMS based machines. Projects within 2600 should continue to follow the development of industry standard protocols. As protocols evolve there must be flexibility on the part of not only 2600 but also the LAN user community to adopt the new standard. If there are changes in the protocol, projects in 2600 would be responsible for providing migration paths (to minimize impact on application interfaces), and modifying/testing the protocol to adhere to security standards/conventions. User modification of the protocol should be restricted and should require approval and testing by 2600 personnel to ensure that security requirements are met.

Network functions

NFS-The Network File System is the de facto industry standard for transparent storage, access, and sharing of files between workstations, i.e. across the network. Although NFS specifications appear to be consistent on the surface, there are some features dependent upon vendor implementation. Having knowledgeable personnel in the support groups is essential. Procedures need to be established for the proper administration of the NFS service within a LAN and across LANs. These procedures should be adaptable to meet the needs and flexibility required for the different LAN user communities and the possibility of establishing a central NFS server within the CCN.

Basic Services should be common across the workstation/LAN environment. Support of common network services such as Telnet, FTP, Mail, etc. should be available.

Remote execution and distributed computing-Apollo's Network Computing System (NCS), which is being promoted by Apollo as a standard, allows inter-vendor distributed computing. Capabilities include a mechanism for calling subroutines across the network as well as the ability to dynamically spread such computing tasks across the network. While such functionality is desirable, it is premature to recommend standardization on a particular interface in this area.

System and Utility Software

Operating systems must be versions which include features that satisfy security requirements specified by 2600. It is possible that 2600 may need to make local security modifications to meet those requirements, or in conjunction with a vendor, assure the necessary functionality is included in vendor supported software. Any modifications to the operating system must be accomplished or approved and then tested by the responsible support group. Failure to abide by this could cause serious security and/or integrity problems.

With regard to UNIX based machines, the Workstation Project does not mandate a specific standard for UNIX. Either AT&T System V or Berkeley based UNIX systems may be acceptable if they include the necessary standard features to meet requirements. We believe that UNIX will converge to a single core operating system. Indeed there is vendor commitment toward this goal. But, should support and functional requirements determine a need to select a more specific standard, then we would recommend strategies based on System V.

Windowing systems-X-Windows is rapidly becoming the industry standard and should be the windowing system that 2600 currently supports in the development of applications. When tools are developed by 2600 for the workstation/LAN environment and a windowing system is required, the application/tool should be done with X-Windows. If users have a need or requirement for another windowing system this is acceptable, but support, including consulting, for those windowing systems may be low or nonexistent. In general, using a nonstandard windowing system will not affect workstation support in other areas. There are other windowing systems such as the PostScript based NeWS system that could influence future windowing standards. Systems should be supported as they become standard and/or there is a wide-spread user requirement for their support.

Compilers-several compilers are available for workstations-too numerous to list here. The C language seems to be the most prevalent language in use on workstations and thus should be considered for implementation of utilities, tools, libraries, etc. Other languages such as Fortran and Ada will be required and essential to users. Interfaces (especially from such languages as Fortran) to lower level C libraries would be necessary. Consulting support should be available in these areas.

Common Usage Software

Document Preparation-There are several word processors/document preparation software packages available for workstations. The current accepted standard at Sandia-MASS-11-is not directly available on UNIX based machines, but could be used with a DOS emulation window. This document was prepared with Frame Maker, a document publishing package available for Sun (and a few other) workstations. WordMARC is available and we have it installed on our Sun workstations. Another document preparation package is from Interleaf-it is

available for a variety of workstations and is a popular package. We are not establishing any specific recommendations or guidelines for support in this area.

editors-Although there are numerous editors available, *vi* and *ed* are the most prevalent editors for UNIX based machines. Users should expect consulting and/or training for these common editors.

graphics software-Support would be based on the existing and/or emerging standards. These include GKS, PHIGS and PHIGS+, and the Computer Graphics Metafile (CGM).

software control-Both RCS (Revision Control System) and SCCS (Source Code Control System) are widely available on UNIX based machines. In conjunction with these utilities, the *make* utility, which maintains, updates, and regenerates groups of programs, is available to describe dependencies among programs and files, and to define rules when specific actions are required. Consulting and training support should be available.

Other

Certainly there will be other software that could be included in this list, especially if there is a need within the user community. Also, software tools and applications will be developed within 2600 for the system management, graphics, and security areas which may become standardized. Thus, ways for doing a particular task, whether it is adding a new user, checking to see that security features are in place, interfacing with a graphics processor, or whatever, should be consistent and simplified for both the user community and 2600 personnel.

Hardware Standards/Conventions

Hardware Standards and Conventions, like Software Standards and Conventions, are essential in order to maintain control, flexibility, and consistency of the workstation environment. Recommendations and guidelines with respect to communications, workstation vendors, and personal computers follow.

Communications-Hardware Standards and Conventions are essential in the communications area. If a user wishes to connect a workstation to 2600 communications facilities he must be using those devices supported by the communications divisions (2647/48). The PBX Ethernet capability will be a significant interface used to connect workstations/LANs to central site. A fiber optic star configuration is another possibility for the interconnection of workstations. Communication Policy/Standards are addressed in detail in SNL, Albuquerque LAN Communications Policy and SNL, Albuquerque LAN Communications Standards (2647).

Workstation Vendor-The Project has no recommendation for a specific workstation (vendor). The requirements are that the workstation meet the standards as specified within this procedures/Standards/Conventions section. We expect through the adoption of standards/conventions and the future unifying of UNIX that a significant portion of the 2600 support will be applicable to most if not all workstations. It is not recommended that 2600 increase manpower just to have knowledge of all vendor workstations, but we would expect manpower increases in certain areas as specified in the Support Group section.

Personal Computers-Within the market place the difference between a high end PC and a low end workstation is somewhat indistinguishable. Indeed with the proliferation of PCs already at Sandia-adaptation of these with third party boards and software is inevitable. Thus, there are and will be several PCs that fit into the standards specified and should be considered for support-no different than a workstation.

Implications of Adopting Standards/Conventions on the CCN

Adoption of standards/conventions in the workstation environment will tend to drive a need for the same or similar standards/conventions for the CCN environment, especially with respect to distributing the computing environment in a client/server flavor. The question then arises: "Will the CCN support the functionality of the LAN environment?" The Network Architecture Project is looking into standards/conventions for the network and the Prototype Secure Network Project should investigate the implication and possible acceptance of this environment. The LAN environment provides enhanced functionality over existing capability.

Network services that should be considered by 2600 include but are not limited to:

- o NFS support on central machines
- o X-Windows running on central computing systems that have user login capabilities
- o Remote logins (to those machines accepting user logins)
- o direct file transfer between machines (similar to FTP)
- o NFS server in the CCN
- o Distributed network computing through remote execution capabilities such as NCS
- o Open UNIX mail facility
- o Open central UNIX computer

Because most of the above services are already in place for TCP/IP environments and due to the difficulty and increased manpower that would be required to incorporate

some of the above into our existing CCN, we would recommend commitment toward TCP/IP facilities in the short term. Even as movement is made toward TCP/IP, investigation into the first five services should be underway. The long term recommendations would be to go in the direction of emerging ISO standard protocols. Taking the first step toward TCP/IP should make this long term goal easier to accomplish.

Appendix B.

Glossary of Common SNL, Albuquerque Networking Terminology and Acronyms

ACL Access Control List

For VAX computers, for example, a list that defines the kinds of access to be granted or denied to users of 'objects' such as files, devices, and mailboxes.

ACM Association for Computing Machinery

An organization for computer professionals concerned primarily with computer software.

ADNET ADministrative NETwork

A system of distributed IBM mainframe nodes that process administrative data.

AFT A File Transfer

A LANL-developed file transfer protocol. It is used by MASS.

AIS (1) Administrative Information Systems

Several applications in various Sandia departments which administer data along functional lines. Some major systems are:

ACS	Access Control System
FIS	Financial Info. Systems
HRIS	Human Resources Info. System
ICS	Inventory Control System

AIS (2) Administrative Information Systems

IPARS	Integrated Procurement System, Accounts Payable, & Receiving Systems
IPS	Integrated Procurement System
MIDAS	Management Info. Data Access System
PIMS	Purchasing Info. System
PRS	Personnel Reporting System
PRS/PC	A PC application of PRS

AIS (3) Administrative Information Systems

REGAL	Report Generating Algorithm
REGENT	Report Generator Terminal Reporting System
TERSE	Terminal Reporting System

ANSI American National Standards Institute

An association that develops American industry standards for coded character sets, data transmission, programming languages, computer definitions, magnetic ink, optical character recognition, etc.

ARPA Internet Advanced Research Projects Agency Inter-network Communications

An internetwork of several networks all running the TCP/IP protocol suite. It connects research agencies, commercial sites, military sites, and universities.

ARPANET Advanced Research Projects Agency NETwork

A subnet of the ARPA Internet. Agencies connected to the ARPANET are all federally funded.

ASCII American Standard Code for Information Interchange

The 128-character set used in telecommunications.

ATL Automatic Tape Library

A hardware component of the Integrated File Store that contains magnetic tapes - generally used to store very large datasets or data which are used infrequently.

AVS Access Validation System

The system that verifies passwords and/or user IDs for SNL, Albuquerque's secure computing partition.

C C

A programming language originally developed at Bell Laboratories for creating the UNIX operating system. It has become the language of choice for many systems applications developed on that system and elsewhere.

CCF Central Computing Facility

(Also called 'The Annex.') The physical facility containing the scientific and administrative large mainframe and support computers operated by Org. 2600 for general Laboratory use. The Central Computing Network (CCN) is housed in the CCF.

CCN Central Computing Network

The current SNL, Albuquerque network composed of the HYPERchannel and several worker/service scientific and administrative computers (like the Cray and IFS) housed in the Central Computing Facility.

CFS Common File Store

The software package (developed for DOE) that performs security and data management for the Integrated File Store system.

CGI Computer Graphics Interface

A proposed standard for the interface to a virtual graphics device.

CGM Computer Graphics Metafile

A standard for file content and format which is used for storage and retrieval of picture information.

COM Computer Output Microfilm/fiche

A term commonly used to mean output which has been printed (or plotted) on some type of photographic film ranging from 16mm up through 105mm microfiche.

COMET COMET

A stand-alone, dialup or PBX(TSN)-accessible electronic mail system available to Sandians. It also is used by contractors and DOE agency personnel around the country to communicate with Sandians.

CPU Central Processing Unit

The hardware that handles all calculating and routing of input and output as well as executing programs. In short, the CPU is the part of the computer that actually computes.

CRAY Cray Research, Inc.

A manufacturer of large mainframe computer systems, several of which are used at Sandia (e.g., Cray X-MP, Cray 1S).

CSMA/CD Carrier Send Multiple Access with Collision Detection

A local area network access method. This format is used by Ethernet.

CTSS Cray TimeSharing System

An interactive operating system, developed originally at Lawrence Livermore Laboratories to provide timesharing service for users of the Cray computer. Various versions have been in use at SNL, Albuquerque since 1985.

DNTS Distributed Network Terminal Service

Software which runs on the HTS machines that allows users to request interactive sessions on the Cray computers either from their local VAX or through the Terminal Switching Network.

DOENTS Dept. of Energy National Telecommunications Services

A nationwide communications network operated by DOE for the DOE and its contractors. There is a node at SNL, Albuquerque for access to DOENTS. Access is available to SNL, Livermore, LLNL, Nevada test sites, ALO, and many other sites.

DQS Distributed Queue Service

DEC software which allows setup/use of remote printers as if they were physically located on the user's local VAX.

EBCDIC Extended Binary Coded Decimal Interchange Code

An eight-bit code used to represent 256 unique alphabetic characters, numbers, and special characters.

ELSENET Extended Laboratory SErvice NETwork

A network which can be accessed by secure VAX Systems on the SNL, Albuquerque VAX Network, but not accessible via the TSN. ELSENET consists of nodes at KCD, SNL, Livermore, GEND, and SNL, Albuquerque. Classified data can be sent only to the KCD node currently.

Ethernet Ethernet

A local area network protocol which uses CSMA/CD (Carrier Send Multiple Access with Collision Detection).

FCE Future Computing Environment

A completed project which helped to define the 'future computing needs' of users.

FDDI Fiber Distributed Data Interface

An emerging 100 Mb/s fiber-optic based communications standard.

FOCUS Facility for Operations Control & Utilization Statistics

The software package running on the OPM computer which is performing the CTSS batch job scheduling function via the PROD software.

FORTRAN FORmula TRANslator

A programming language for writing programs in a mathematical style language which is then translated into a machine language for execution. FORTRAN is the most widely used programming language for developing programs dealing with scientific applications.

FTP File Transfer Protocol

A commonly available TCP/IP application which permits transfer of files between nodes on a network.

GCS Graphics Compatibility System

This public domain, general purpose, graphics subroutine library offers 2-Dimensional and 3-Dimensional drawing capabilities, plotting, and graphical fonts currently in use at SNL, Albuquerque.

GDDM Graphical Data Display Manager

IBM's primary (mainframe) device support and host graphics software used to create presentations of text, image, and graphical data on terminals, printers, and plotters.

GKS Graphics Kernal System

A standard interface for doing low level, 2-Dimensional graphics. It has segmentation capabilities and supports simultaneous multiple workstations. It has limited support at SNL, Albuquerque.

GOSIP Government Open System Interconnect Profile

GOSIP specifies a set of OSI protocols for computer networking that are intended for acquisition and use by government agencies.

HSX High Speed eXternal communications channel

A high speed external communications channel used on the Cray X-MP (100 Mbytes/sec transfer rate) and Cray-2. It has the ability to communicate with certain non-Cray equipment as well as providing high speed communications between Crays.

HTS Hyperchannel Terminal Server

These machines allow asynchronous terminals access to the worker computers on the CCN HYPERchannel backbone. Users entering via the HTS machines are validated by the AVS and may then log into CTSS.

HYPERCHANNEL HYPERchannel

Network Systems Corporation's high speed intercomputer data connection methodology, consisting of both hardware and software.

IAS Information Access System

Software being developed by Org. 2613 personnel for use on personal and VAX computers. The purpose is to provide easy network access to existing applications and services.

IEEE Institute of Electrical & Electronic Engineers

A major professional organization of engineers which sometimes is instrumental in defining standards for certain computer hardware and software items.

IFS Integrated File Store

The central file storage system serving all CCN equipment and some distributed systems. In order to supply security and data management, a DOE-sponsored software package called CFS runs with IFS software to perform file management.

IOP Input Output Processor

An individual processor of a Cray X-MP Input Output Subsystem (IOS).

IOS Input Output Subsystem

A Cray X-MP hardware subsystem that performs the input/output functions. An IOS includes 2 to 4 IOPs.

ISEC Insider SECurity program

An internal SNL security program designed to identify patterns of performance or activities that could forewarn of security problems.

ISO/OSI International Standards Organization/Open Systems Interconnect

An organization dedicated to the propagation of standards in the international computing/networking environment.

K K

A unit for measuring the size of memory or similar resources. K is short for kilo, and is used roughly to mean 1000, although K is equal to 2^{10}, or 1024.

KCD Allied Signal, Kansas City Division

KERMIT KERMIT (named after the famous frog)

A terminal emulator and file transfer protocol available for a very wide variety of microcomputers and mainframes. (Developed at many sites and distributed by Columbia University)

LAN Local Area Network

A network of computers (often PC's, workstations, and/or minicomputers) with an intelligent communications protocol like Ethernet or token ring. The connected computers are usually in a restricted geographic area.

LANL Los Alamos National Laboratory

LLNL Lawrence Livermore National Laboratory

MAKE Make

A standard UNIX system utility which will execute a series of stored parameterized commands based on the existence and date/time stamp of files.

MASS MASS

The user interface utility program on a computer connected to the CCN which allows access to the IFS from that computer.

MASSNET MASSNET

A protocol which runs on the HYPERchannel and attached computers. It is used in applications such as MASS, DNTS, and PROD.

Mb Megabit

One million data bits.

MB Megabyte

One million data bytes.

MILNET MILitary NETwork

A subnet of the ARPA Internet which connects military-related facilities.

MODEM MOdulator DEModulator

A device which translates information generated by a computer into signals which can be transmitted over phone or hardwired lines to another location, where another modem translates the signals back into computer-understandable information.

NA Network Adaptor

Hardware required to interchange protocols from a specific computer to the Massnet protocol on the HYPERchannel.

NCS Network Computing System

Software developed by Apollo Computer Inc. which provides distributed, inter-vendor computing at the subroutine level in a heterogeneous network environment.

NETEX NETwork EXecutive

A proprietary communications protocol from Network Systems Corp. used on the HYPERchannel. It is used in addition to MASSNET on the CCN HYPERchannel.

NeWS Network-extensible Windowing System

A network window system, developed by Sun Microsystems, Inc., which uses PostScript together with Sun extensions as a windowing protocol.

NFS Network File System

The de facto industry standard developed by Sun Microsystems, Inc. for transparent storage, access, and sharing of files between dissimilar computers across a TCP/IP network. Sharing is done at the buffer level, without need to copy files from one host to another.

NQS Network Queueing System

A public domain system developed by for the queueing of devices on UNIX systems.

NSC Network Systems Corporation.

The manufacturer of HYPERchannel adapters and associated software. They supplied SNL, Albuquerque's HYPERchannel and some of the software used with it.

NWCNET Nuclear Weapon Complex NETwork

A network on which files are moved between CFS storage systems at different DOE laboratories. The WBCN is used in this process.

ONODE Output NODE

A network node which performs printing and plotting (to create paper, microfiche, other film, videotape) at the Central Site for CCN mainframes as well as distributed computers. The user interface for the Output Node is also called ONODE.

OPM Operations Performance Monitor

A machine used to submit all batch jobs to the Cray computers running CTSS. It schedules batch jobs using the FOCUS software.

OPMODEL OPerations MODEL

A DOE satellite based, four-node, intercity communications network. It has both voice and data capabilities with transmission rates up to 1.5 MB/s. The network currently connects DOE/HQ, Allied Signal/KC, SNL, Albuquerque, and LNLL.

OS/2 Operating System/2

A new operating system for 386-based personal computers.

OSI Open System Interconnection

An abstract 7-layer model devised by the International Standards Organization (ISO) to describe communications between computers. OSI sometimes refers to the protocols being developed by the ISO (within the context of the 7-layer reference model).

PACX Private Automatic Computer eXchange

A old data contention and switching system to maximize utilization of computer resources through port switching and contention. Different PACX systems served different physical areas. (Also called Port Contender or Data Switch.)

PBX Private Branch eXchange

A new secure data switch which replaced the PACX. It provides asynchronous, synchronous, Ethernet and IBM 3470 services.

PC Personal Computer

One of a series of 'new generation' computers that reside on the user's desk.

PHIGS Programmer's Hierarchical Interactive Graphics System

A proposed standard for creating low level 2-Dimensional and 3-Dimensional graphics.

PON Protected Open Network

The new physical network being built at SNL, Albuquerque to support the open scientific information partition.

POSIX POrtable Standard unIX operating system

A proposed standard for UNIX systems. It currently, only defines the interface which the operating system provides to "C" programs.

PROD PRODuction

The software package that runs on the OPM computer as a subfunction of FOCUS and on the CTSS Cray computers and VAX computers to invoke batch job submission.

PSN Prototype Secure Network

A new physical network being built at SNL, Albuquerque to develop, implement, and test the elements of the SSN. It will also be used to test and demonstrate UNIX functionality.

RCS Revision Control System

A public domain source code management utility which is commonly available for UNIX environments.

RLOGIN A UNIX facility for doing remote login on TCP/IP networks.

A public domain source code management utility which is commonly available for UNIX environments.

SCN(1) Sandia Computing Newsletter

A universe-famous monthly publication of Org. 2614 which contains articles of interest to the Sandia computing community.

SCN(2) Secure Computing Network

SNL, Livermore's secure network.

SCSC Small Computer Support Center

A capability to help users in choosing PC hardware and software. New products are evaluated and, where of general interest, made available to Sandians.

SDM/S System Development Methodology/Structured

The system life cycle method for development of software, with steps and forms for using the methodology to ensure that all phases are planned and completed.

SLATEC Sandia/Los Alamos/AFWL Technical Exchange Committee

A technical committee made up of representatives from the three laboratories to exchange technical computing information and to jointly develop common software packages in the areas of graphics and mathematical libraries. (AFWL stands for Air Force Weapons Laboratory)

SNA Systems Network Architecture

The IBM network which will be used in support of the ADNET system to connect together the area IBM administrative computers.

SSD Solid-State storage Device

A semiconductor memory subsystem that acts as a very fast access secondary storage device for the Cray X-MP.

SSN Secure Supercomputer Network

The new physical network being implemented to support the secure scientific information partition at SNL, Albuquerque.

SVDI Sandia Virtual Device Interface

A locally developed interface for supporting device-independent, graphical input/output to a variety of graphics devices. Functionality is basically 2-Dimensional vector, polygon, minimal text and input. (also called VDI)

SVDI + R SVDI + Raster

An extension of SVDI which includes functions for raster output (supports output of pixel data). (also called VDR)

TCC Technical Control Center

Locations within SNL, Albuquerque where PBXs are located along with other communications equipment. TCCs are staffed with technicians who respond to data communications trouble calls. Center technicians also maintain communications equipment & implement DCR requests.

TCP/IP Transmission Control Protocol/Internet Protocol

A commonly used suite of communications protocols used in ARPA Internet, some UNIX systems, and others. This term is sometimes used to refer to the entire DARPA protocol suite.

TELNET Telnet

The ARPA service facility for doing remote login on TCP/IP networks.

TSN Terminal Switching Network

An analog communications network of wire and fiber optics providing terminal communications throughout SNL, Albuquerque. The network contains data switches (currently PBXs) to multiplex data among buildings, allowing any given terminal to connect to a variety of host computers, although not simultaneously.

TSN-Secure Terminal Switching Network (Secure)

The TSN used for local terminal communications can also be used to gain access to GEND; SNL, Livermore; KCD; Tonopah; and LANL. Special classes of service exist that provide for access to these sites.

TTR Tonopah Test Range

UNICOS UNICOS

Cray's version of the UNIX operating system.

UNIX UNIX

An operating system developed at Bell Labs. It is gaining in popularity as the operating system of choice for many computer systems.

VI vi

The standard full-screen visual editor developed at Berkeley as an enhancement to the "ex" editor, and provided with most UNIX systems.

VLSI Very Large Scale Integration

A computer chip technology which is able to place in excess of 100,000 gates on a single chip.

WAN Wide Area Network

Any network larger than a local area network. Some consider a WAN to be city, state, national, or international-wide.

WBCN Wide Band Communications Network

This system supports classified data transmission between production facilities and the DOE weapons labs (SNL, Albuquerque; SNL, Livermore; LANL; LLNL).

Windows Windows

A software system which allows a user and/or application to manage multiple, simultaneous activities on the same display, with each activity operating in its own "window". Each active window generally occupies some area of the screen, although at times a window may be obscured by other windows or temporarily closed.

X-Windows X-Windows

A public domain windowing system developed by Massachusetts Institute of Technology which is emerging as a de facto industry standard.

X-MP X-Multi-Processor

A particular model of Cray supercomputer which features 2 or 4 processors, a solid state disk, and an input/output subsystem.

XENIX XENIX

A PC implementation of UNIX developed by Microsoft Corporation..

XNET XNET

The file transfer software which executes on the SNL, Albuquerque's VAX Gateway connecting the CCN and the Secure Distributed VAX Network. It was originally developed at LANL.

DISTRIBUTION:
1100 F. L. Vook
1110 S. T. Picraux
1111 B. L. Doyle
1112 S. M. Myers
1113 R. C. Hughes
1120 J. P. Gerardo
1124 A Owyoung
1126 A. W. Johnson
1128 G. N. Hays
1130 G. A. Samara
1131 B. Morosin
1134 T. A. Michalske
1140 P. S. Peercy
1141 H. T. Weaver
1143 P. L. Gourley
1150 J. E. Schirber
1151 D. R. Jennison
1152 D. W. Schafer
1153 B. Morosin
1164 A. Owyoung
1200 J. P. VanDevender
1220 M. Cowan
1221 M. M. Widner
1230 J. E. Powell
1231 J. E. Lee
1232 W. Beezhold
1233 L. M. Choate
1234 R. J. Leeper
1235 J. M. Hoffman
1236 G. A. Zawadzkas
1237 G. W. Kuswa
1240 K. R. Prestwich
1241 J. R. Freeman
1242 D. E. Hasti
1245 J. J. Ramirez
1248 M. T. Buttram
1260 D. L. Cook
1261 G. O. Allshouse
1262 D. L. Cook, Actg
1264 R. W. Stinnett
1265 J. P. Quintenz
1266 D. D. Bloomquist
1270 J. K. Rice
1272 D. H. McDaniel
1273 M. K. Matzen
1274 R. J. Lipinski
1275 R. A. Gerber
1400 E. H. Barsis
1410 P. J. Eicker
1411 J. J. Wiczer
1412 P. J. Eicker Actg.
1414 R. W. Harrigan
1415 K. T. Stalker
1420 W. J. Camp
1421 R. J. Thompson
1422 R. C. Allen
1423 E. F. Brickell
1424 G. G. Weigand
1500 W. Herrman
1510 J. W. Nunziato
1511 D. K. Gartling
1512 J. C. Cummings
1513 D. W. Larson
1520 L. W. Davison
1521 R. D. Krieg
1522 R. C. Reuter Jr.
1523 J. H. Biffle
1524 A. K. Miller
1530 D. B. Hayes
1531 S. L. Thomson
1533 P. Yarrington
1534 J. R. Asay
1550 C. W. Peterson
1551 J. C. Cole
1552 D. D. McBride
1553 S. McAlees Jr.
1554 D. P. Aeschliman
1555 W. R. Barton
1556 W. L. Oberkampf
1800 R. J. Eagan
1810 D. W. Schafer
1811 R. L. Clough
1812 J. M. Zeigler
1813 J. G. Curro
1815 G. E. Pike
1820 J. B. Woodard
1821 S. H. Weissman
1822 K. H. Eckelmeyer
1823 J. A. Borders
1830 M. J. Davis
1832 W.B. Jones
1834 A. K. Hays
1840 R.E. Loehman
1842 B. C. Bunker
1846 D. H. Doughty
2000 G. T. Cheney
2100 H. J. Saxton
2110 R. E. Bair
2113 B. D. Shafer
2114 M. G. Knoll
2115 F. W. Hewlett
2130 D. L. Weaver
2131 R. W. Light
2132 P. Plunkett
2140 H. J. Saxton, Actg
2142 R. E. Anderson
2146 T. A. Dellin
2147 P. S. Winkur
2170 E. D. Graham
2300 J. L. Wirth
2310 M. K. Parsons
2311 T. D. Donham
2312 D. J. Allen
2314 H. L. Floyd Jr.
2320 J. H. Renken
2321 E F. Hartman
2322 G. J. Scrivner
2330 J. H. Stichman
2331 A. E. Sill
2334 G. M. Heck
2335 H. M. Witek
2336 E. J. Nava
2340 M. W. Callahan
2341 C. R. Alls
2342 D. W Arquette
2343 W. H. Schaedla
2344 C. S. Williams
2345 B. C. Walker
2346 R. L. O'Nan
2360 P. A. Longmire
2361 T. J. Williams
2362 M. B. Murphy
2363 H. M. Barnett
2364 J. F. Jones Jr.
2500 R. L. Schwoebel
2510 D. H. Anderson
2512 J. G. Harlan
2513 D. E. Mitchell
2514 L. L. Bonzon
2515 P. D. Wilcox
2520 N. J. Magnani
2525 R. B. Diegle
2526 R. C. Lincoln
2530 G. M. Ferguson, Actg
2533 G. L. Laguna, Actg
2534 H. M. Bivens
2540 G. N. Beeler
2541 J. P. Abbin
2542 W. N. Sullivan
2543 K. G. McCaughey
2545 J. H. Barnette
2560 J. T. Cutchen

2561 R. A. Damerow
2564 F. M. Bacon
2565 J. M. Wilder
2600 L. D Bertholf
2601 J. C. Garrison
2601 F. Cheek-Martin
2601 M. B. Marrs
2601 B. Williams
2601 Staff (11)
2610 D. C. Jones
2612 B. Stiefeld
2612 Staff (13)
2613 L. D. Daigle
2613 Staff (10)
2614 A. R. Iacoletti
2614 Staff (11)
2616 S. A. Love
2616 Staff (4)
2620 K. O. Waibel
2621 G. L. Esch
2621 Staff (8)
2624 T. L. Ferguson
2624 Staff (14)
2625 P. L. Murphy
2626 D. S. Rarick
2629 J. A. Ruggles
2630 W. F. Mason
2631 R. M. James
2631-1 J. E. Gonzales
2631-1 Staff (8)
2631-2 A. R. Gibson
2631-3 M. Foster
2631-4 D. McGinnis
2632 G. E. Connor
2632-1 B. R. F. Collins
2632-2 G. O. Concannon
2632-3 M. C. Finley
2632-5 T. R. Garcia
2634 S. K. Fletcher
2634 Staff (16)
2635 R. E. Jones
2635 Staff (13)
2640 E. J. Theriot
2640A L. Stans(5)
2643 L. D. Buxton
2643 C. M. Schaldach(5)
2643 Staff (8)
2645 W. D. Swartz
2645 C. D. Brown(5)
2645 S. M. Kelly(5)
2645 Staff (7)
2647 M. O. Vahle
2647 L. G. Pierson(5)
2647 Staff (8)
2648 D. M. Darsey
2648 Staff (7)
2800 W. L. Alzheimer
2810 D. W. Doak
2811 T. R. Perea
2812 J. Q. Wisniewski
2813 J. R. Yoder
2814 P. A. Erickson
2820 Gino Carli
2821 G. F. Quinlan
2825 J. K. Sharp
2826 A. J. Ahr
2830 G. R. Urish
2831 T. A. Lovato
2832 L. D. Byers
2833 S. S. Baca
2834 K. E. Wiegandt
2850 D. L. McCoy
2851 L. K. Grube
2852 D. J. Klinetobe
2853 J. A. Lovato
2854 R. E. Thompson
2855 D. L. McCoy, Actg
2857 W. B. Bopp
2858 W. Drozdick
3141 (5)
3151 (3)
3141-1 (8) C. L. Ward
for DOE/TIC
3212 G. E. Tucker
5100 H. W. Schmitt
5110 C. C. Burks
5111 D. N. Bray
5112 L. J. Merrell
5113 R. W. Jorgensen
5114 W. D. Ulrich
5115 S. D. Meyer
5120 W. R. Reynolds
5122 J. N. Baker
5128 B. E. Bader
5140 C. M. Tapp
5141 C. A. Harris
5142 C. M. Tapp, Actg
5143 R. D. Robinett
5144 D. E. Ryerson
5150 K. D. Hardin
5151 W. C. Nickell
5160 G. R. Otey
5161 K. D. Nokes
5165 J. M. Freeman
5166 R. C. Hartwig
5174 P. D. Merilat
5200 J. Jacobs
5210 C.C. Hartwigsen
5212 J. C. Cutter
5213 W. N. Sullivan
5214 W. F. Hartman
5217 D. L. Mangan
5219 R. W. Moya
5230 M. L. Kramm
5231 E. R. Julius
5234 C. L. Schuster
5240 D. S. Miyoshi
5249 B. J. Steele
5250 T. A. Sellers
5251 E. A. Chipman
5253 P. D. Merillat
5253 M. K. Bencoe
5256 R. F. Glaser
5260 J. R. Kelsey
5267 S. C. Roehrig
5268 S. J. Weissman
6010 D. E. Arvizu
6012 B. H. Van Domelen
6200 V. L. Dugan
6216 J. T. Holmes
6217 P. C. Klimas
6220 D. G. Schuler
6221 E. C. Boes
6223 G. J. Jones
6224 D. G. Schueler, Actg
6225 H. M. Dodd
6230 W. C. Luth
6231 H. C. Hardee
6233 T. M. Gerlach
6250 R. K. Trager
6252 J. C. Dunn
6253 D. A. Northrop
6257 J. K. Linn
6258 P. J. Hommert
6300 R. W. Lynch
6310 T. O. Hunter
6311 A. L. Stevens
6312 F. W. Bingham
6313 T. E. Blejwas
6314 J. R. Tillerson
6315 L. E. Shephard
6320 J. E. Stiegler
6321 R. E. Luna
6322 G. C. Allen Jr.
6330 W. D. Weart
6332 L. D. Tyler

6334 D. R. Anderson
6400 D. J. McCloskey
6410 D. J. McCloskey, Actg
6412 A. L. Camp
6413 E. D. Gorhman-Bergeron
6415 R. M. Cranwell
6416 E. J. Bonano
6418 J. E. Kelly
6420 J. V. Walker
6421 P. S. Pickard
6423 K. O. Reil
6425 S. S. Dosanjh
6427 M. Berman
6429 K. D. Bergeron
6440 D. A. Dahlgren
6442 W. A. vonRiesemann
6447 M. P. Bohn
6450 T. R. Schmidt
6451 T. F. Luera
6452 J. W. Bryson
6453 J. S. Philbin
6454 G. L. Cano
6500 A. W. Snyder
6510 W. B. Gauster
6511 L. O. Cropp
6512 F. E. Haskin
7100 C. D. Broyles
7110 J. D. Plimpton
7111 J. R. Banister
7112 J. D. Plimton
7116 S. R. Dolce
7120 M. J. Navratil
7124 P. L. Nelson
7125 R. L. Rutter
7126 C. W. Cook
7130 J. D. Kennedy
7132 A. B. Church
7133 O. L. Burchett
7170 R. D. Bentley
7171 R. R. Beasley
7172 R. G. Hay
7173 G. L. West
7200 C. H. Mauney
7210 R. E. Sonnenberg
7213 J. W. Mead
7220 R. R. Prairie
7222 G. C. Novotny
7223 R. G. Easterling
7230 J. F. Ney
7231 H. D. Bickelman
7232 S. D. Spray
7233 R. E. Smith
7234 J. M. Portlock
7240 R. T. Johnson
7241 L. J. Azevedo
7242 B. E. Barnaby
7250 G. T. Merren
7251 J. M. Sjulin
7252 D. W. Bushmire
7253 R. E. Baack
7254 M. A. Blackledge
7260 J. A. Hood
7261 J. L. Duncan
7262 F. A. Ross
7263 J. N. Middleton
7265 T. Starr
7266 F. W. Muller
7400 J. C. King
7470 J. L. Ledman
7472 J. A. Sayre
7473 T. P. Conlon, Jr.
7476 G. L. Cessac
7480 J. K. Walter, Jr.
7483 P. W. Plomp
7484 W. R. Leuenberger
7500 D. M. Olson
7520 T. J. Hoban
7522 F. D. Gutierrez
7525 R. L. Eno
7526 L. W. Lathrop
7530 T. L. Workman
7533 F. H. Mathews
7535 D. C. Bickel
7540 T. B. Lane
7541 T. J. Baca
7542 T. G. Priddy
7543 R. Rodeman
7550 T. S. Edrington
7551 O. J. Burchett
7553 M. V. Morris
7554 L. D. Scott
7555 B. Boverie
7556 R. A. Hill
7800 R. W. Hunnicutt
8024 (1)
8100 E. E. Ives
8130 L. A. Hiles
8131 W. R. Bolton
8132 A. S. Rivenes
8133 A. C. Skinrood
8150 J. B. Wright
8154 D. J. Havlik
8155 J. B. Woodard
8156 J. B. Wright, Actg
8160 D. J. Bohrer
8161 M. H. Reynolds
8162 C. T. Yokomizo
8163 W. G. Wilson
8164 P. G. Heppner
8165 G. C. Story
8170 J. W. Hickman
8171 C. B. Layne
8172 R. J. Gallagher
8173 M. E. John
8174 O. Schreiber
8200 R. J. Detry
8230 W. D. Wilson
8231 M. H. Pendley
8233 J. M. Harris
8234 C. D. Skogg
8235 D. L. Crawford
8236 P. W. Dean
8240 C. W. Robinson
8241 G. A. Benedetti
8242 M. R. Birnbaum
8243 M. L. Callabresi
8244 C. M. Hartwig
8245 R. J. Kee
8270 R. C. Dougherty
8274 J. E. Marion
8280 R. A. Baroody
8281 J. D. Benton
8300 P. L. Mattern
8310 R. W. Rohde
8312 J. A. Brooks
8313 D. L. Lindner
8314 M. W. Perra
8316 J. B. Woodard
8317 A. R. Willis
8340 W. Bauer
8341 M. I. Baskes
8342 M. Lapp
8343 R. H. Stulen
8347 K. Wilson
8350 J. S. Binkley
8351 R. P. Lucht
8353 G. A. Fisk
8354 R. A. Palmer
8357 R. W. Carling
8360 W. J. McLean
8363 B. R. Sanders
8364 S. C. Johnston
8400 R. C. Wayne
8430 J. Vitko
8431 L. A. West

8432 M. Hall
8435 D. W. Sweeney
8436 J. Vitko, Actg
8440 H. Hanser
8441 A. T. Jones
8442 E. English
8445 E. T.Cull
8446 A. J. West
9000 R. L. Hagengruber
9010 W. C. Hines
9011 G. S. Brown
9012 C. C. Carson
9013 W. H. Ling
9014 J. W. Keizur
9015 R. R. Preston, Jr.
9100 R. G. Clem
9110 C. W. Childers
9111 R. L. Ewing
9112 S. A. Dupree
9113 C. S. Lee
9114 B. F. Johnson
9115 J. M. Fletcher
9120 M. M. Newsom
9121 L. T. James
9122 R. H. Braasch
9127 C. E. Olson
9130 R. D. Andreas
9131 D. D. Boozer
9132 A. C. Watts
9133 L. D. Hostetler
9140 D. J. Rigali
9141 M. W. Sterk
9142 A. C. Bustamante
9143 E. W. Reece
9144 W. E. Williamson, Jr.
9200 W. C. Myre
9210 H. M. Dumas
9211 T. G. Taylor
9212 L. W. Maschoff
9213 D. E. Barnes
9214 L. E. Anderson
9220 G. H. Mauth
9221 J. L. Williams
9222 L. S. Walker
9223 D.A. Reynolds
9224 L. J. Ellis
9225 C. A. Boye
9226 W. J. Slosarik
9230 R. Spalding
9231 J. C. Chavez
9233 W. B. Goldrick
9234 R. L. Hostetler
9235 D. E. Miller
9240 P. A. Stokes
9241 J. M. Holovka
9242 K. Fraley
9243 P. B. Herrington