NISTIR 4734

Foundations of a Security Policy for Use of the National Research and Educational Network

Arthur E. Oldehoeft

Chairman
Computer Science Department
Iowa State University

for the

U.S. DEPARTMENT OF COMMERCE
Technology Administration
National Institute of Standards
and Technology
Computer Systems laboratory
Computer Security Division
Gaithersburg, MD 20899

February 1992

U.S. DEPARTMENT OF COMMERCE
Rockwell A. Schnabel, Acting Secretary

NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY
John W. Lyons, Director

Foreword

Security has become a topic of national importance as more and more information is being processed in distributed computer systems and networks. This report explores some underlying considerations in the development of a security policy for the evolving National Research and Education Network (NREN).

The National Institute of Standards and Technology (NIST) is responsible for developing standards and guidelines for the protection of unclassified but sensitive information processed by Federal organizations. NIST standards include technical methods for providing security in cost-effective, interoperable ways. NIST guidelines outline management responsibilities and procedures for effectively and securely using and operating the computers. However, it is the responsibility of each Federal organization to establish its own security policy or policies, to identify and justify the security needed and to establish a security program for implementing and managing the security mechanisms selected.

The National Research and Education Network is one part of a Federal program establishing a comprehensive set of computing and communications services throughout the nation's scientific and educational communities. It consists of a large number of Federal, State and commercial entities providing and using a wide range of services. As such, it will be organized as a cooperative with a membership including service suppliers and users with a distributed management structure. Basic security services will be expected by most users with special security services required by others.

Specifying security policies, procedures, methods and mechanisms in a loosely coupled, multi-faceted, distributed network such as the NREN will be difficult. Scientific research and education professionals have become accustomed to open access to many computers and data bases without restriction. Simultaneously, they have expected reliable services with a high level of availability, data with a high degree of integrity and communication with some implied level of confidentiality. In practice, there are no assurances of any of these security services and there presently exists no written policy which could be used as a basis for providing them. Future networks must provide easy access to information and information processing services to all those who are authorized in a simple, user friendly manner. Simultaneously, future networks must count on the cooperation of users in order to assure a reasonable level of integrity and availability of processing services and integrity, availability and confidentiality of information to properly authorized users.

This report is the result of a research activity sponsored by NIST. Dr. Arthur Oldehoeft is the Chairman of the Computer Science Department at Iowa State University. During a six month sabbatical leave from Iowa State, he worked at NIST investigating the status of NREN and exploring alternative foundations of a security policy for the NREN. He has analyzed existing security policies and codes of ethics that have been established in several government organizations and university environments. He has coordinated with several leaders in network technology in investigating security policies and provisions that could be acceptable to network implementors, users and managers. However, it has not been reviewed by a large number of users and is not endorsed by any organization having authority over such a network.

This report is the result of a research activity and should not be interpreted as a NIST standard

or guideline established through the Federal Information Processing Standards process. The report is intended for discussion purposes by the people and organizations sponsoring the development and use of the NREN. It will be provided as input to the Federal Networking Council for its deliberations on what security provisions should be expected and provided in the Interagency Interim National Research and Education Network (IINREN). Other security policies and provisions should be expected as this evolves to national and international information networks in the future.

The draft policy outlined in the report is considered a level 1 policy in a theoretical hierarchy of policies. This hierarchy ranges from information technology codes of ethics through a number of refined levels of policies down to implementation specifications for security mechanisms, procedures and protocols that support and enforce the policies. Future research is planned to explore this range of policies in a number of security domains that can be defined within a broad distributed network.

Dennis K. Branstad, Ph.D.
NIST Fellow
Project Advisor

Executive Summary

The National Research and Education Network (NREN) is an integral part of the planned High-Performance Computing and Communication (HPCC) infrastructure that will extend throughout the scientific, technical and education communities. The projected vision is one of desks and laboratory benches as entry points to a nation-wide electronic network of information technologies with shared access to services and resources such as high-performance computing systems, specialized software tools, databases, scientific instruments, digital libraries, and other research facilities.

The problem of computer and network information security (one of the major computing issues of the day), will be complicated by the diversity of requirements as the NREN is designed, developed, and operated in collaboration with potential users in government, industry, research laboratories, and educational institutions. One major impediment to improved security is the lack of a clearly stated security policy for general computing. In recognition of this problem, national organizations are beginning to develop and publish codes of ethics for the use of computers. An Internet working group has recently published guidelines for the secure operation of the Internet.

Recent Congressional legislation for HPCC reaffirms the role of the National Institute of Standards and Technology as the agency that is "responsible for developing and proposing standards and guidelines needed for cost-effective security and privacy of sensitive information in Federal computer systems."

The purpose of this report is to explore the foundations of a security policy and propose a security policy for the NREN, one that is applicable to and identifies responsibilities of all major network constituents: end users, system administrators, management at all levels, vendors, system developers, service providers, and the Federal Networking Council.

In order to establish an appropriate context for the development of a national network security policy and also provide for an understanding of the culture of open computer networks, this report first traces the evolution of "national" networks in the U.S. From the structure and operation of the existing NSFNET and Internet, the probable characteristics of the evolving NREN are projected. Foundations for specification of a policy are established through a review of the basic concepts of "security" and "security policy" and through the examination of existing policies, codes of ethics, and Federal legislation regarding computer information security. A draft policy is then abstractly stated, one that is independent of current technologies and organization-specific practices. Since the development of a widely-accepted and meaningful security policy requires the participation of all major constituents, this draft policy is intended to provide the basis for continuing discussion and further development.

Acknowledgements

The author is pleased to acknowledge helpful discussions with many individuals including Robert Aiken, Dennis Branstad, Vinton Cerf, Steve Crocker, Robert Kahn, Stuart Katzke, Jerry Linn, Robert Rosenthal, and Steve Squires.

Contents

1 Introduction

1.1 Background

After more than two decades of research, security in computer information systems remains one of the "major issues" of the day. While significant technological advances have been (and continue to be) made, the general problem of security has assumed larger and more complex dimensions – attributed in a large part to the development and pervasive use of computer networks.

Computers in industry, universities, and government laboratories are now commonly part of local area networks which are in turn connected to larger regional, national and international networks. Many computers in businesses and homes have dial-up accessibility to other computers in this same web of interconnected networks.

Major efforts have focused on the development of network communication protocols for efficient and reliable transmission of information while less attention has been devoted to the problems of security. In response to a universal acknowledgement of the need for information security, circumstances are changing and increasing attention is being given to security. This need for security is counterbalanced by an equal need to provide functionality for information sharing, distributed computation, file transfer, and electronic mail.

Most of the current research deals with the technical issues of enforcing specific aspects of security - e.g. access controls and encryption. Comparatively little attention has been given to the development of a meaningful and generally accepted "policy" for information security that would be applicable to a diverse set of users and computers in interconnected networks.

The U.S. Congress has recently passed legislation that defines important new initiatives in High Performance Computing and Communications (HPCC). An integral component of the HPCC program is the development of a high-speed "National Research and Educational Network (NREN)" that is intended to link together research and educational institutions, libraries, government laboratories, and industry.[1] As the NREN continues to evolve, the problem of computer information security is expected to become more acute.

1.2 Purpose and Scope of Report

The purpose of this report is to explore the foundations of a national network security policy and propose a draft policy for the NREN. Within this context, a network is defined to be a collection of autonomous computers interconnected by some communication media. The term security refers to computer and network information security. References to the NREN are generally intended to include not only the backbone transmission facilities, but also entities connecting to the NREN: local networks, host computers, and end users themselves. To that end, the scope of the proposed

[1] The NREN is envisioned as the evolution of existing networks, including those of the Federal Agency currently being developed as the Interagency Interim National Research and Education Network (IINREN), to a multi-gigabyte/second network that is expected to mature by 1996.

security policy is broad, addressing the responsibilities of users, managers, system administrators, system developers, vendors, network service providers, and a national coordinating body.

The policy presented here is at an abstract "first-level", somewhat higher than typical policies or directives that establish organization-specific requirements and depend on technology-specific security mechanisms. In order to derive actual security practices from the security policy, further refinements are necessary. Properly addressed, these refinements should culminate in the formulation of formal specifications for security services and supporting mechanisms. Such specifications could be used by service providing organizations or entities, in addition to user organizations, for uniformity and interoperability purposes in providing required protection.

The development of a widely-accepted and meaningful policy is normally an evolutionary process that requires the participation of all major constituents that will be affected by the policy. Therefore, the policy specified in this report is intended to establish a basis of discussion among these constituents on what security services could or should be expected in such a network, what overall security objectives should be pursued in the development of such a network, and who should be responsible for what parts of security.

This report is submitted to the National Institute of Standards and Technology (NIST) as fulfillment of the research effort specified above. It is the result of a six month research activity sponsored by NIST. Under the provisions of the Computer Security Act of 1987, NIST is responsible for developing standards and guidelines for security of unclassified information in Federal information systems. NIST also participates in numerous Federal activities coordinating the use of computer systems in Federal computer networks and in national activities coordinating the use of Federal and commercial systems. This report was developed under sponsorship of NIST (Contract 43NANB112737) as part of its research program in computer security and as a contribution to the Federal Networking Council for its work on the Interagency Interim National Research and Educational Network (IINREN). It is not a proposed standard and has not been reviewed or endorsed by any organization having policy authority over the NREN or IINREN.

1.3 Overview of Report

Chapter 2 traces the evolution of the concept of a "national" network, from its origin to present day considerations. Included is a discussion of the involvement of various Federal agencies and committees. Some projections are made about the ultimate nature of the NREN and its topology in terms of administrative structure that would have impact on the feasibility and enforcement of a national network security policy.

Chapter 3 is a discussion of the foundations for a broad security policy. Basic concepts of security and security policy are reviewed. The need for a national policy is established by noting various codes of ethics, congressional acts, and the diversity of existing organizational policies. Included is a summary of the currently proposed guidelines for secure operation of the Internet.

Chapter 4 is a presentation of the proposed security policy for the NREN. It includes a discussion of scope, vulnerabilities and threats, and responsibilities of the various constituents.

2 Evolution of the NREN

2.1 Introduction

The history of data transmission networks in the U.S. is characterized by astonishing growth, massive innovation, and rapid change. Continuing technological developments indicate that the current pattern of growth will not diminish in the foreseeable future. The NREN, as a next major phase in U.S. research and education networking, is an integral part of the current HPCC Initiative.

As stated in [OSTP91], this initiative has three strategic priorities:

- extend U.S. technological leadership in high-performance computing and communications;
- provide wide dissemination and application of the technologies both to speed the pace of innovation and to serve the national economy, national security, education, and global environment; and
- spur gains in U.S. productivity and industrial competitiveness by making high-performance computing and networking technologies an integral part of the design and production process.

The NREN is the basic component in this planned information infrastructure that will extend throughout the scientific, technical, and educational communities. Its creation is dictated by [BELL88, OSTP91]:

1. the mushrooming demand for electronic communication and sharing of information, including the exchange of more comprehensible supercomputer output and high-quality graphical data;
2. the need for a "collaboration technology" to facilitate cooperation between geographically separated researchers;
3. the Federal interest in maintaining U.S. leadership role in science and technology;
4. the requirements for solving the so-called "Grand Challenges" - defined in [CONG91a] to be fundamental problems in science or engineering with economic and scientific impact, whose solutions will require the application of high performance computing resources;[2] and
5. the need to serve many sectors of government, research, and education.

The vision projected is one of desks and laboratory benches as entry points to a complex high-speed electronic network of information technologies, services and resources, providing access to specialized

[2] In technical terminology, a partial list of challenges would include (Cf. [OSTP91, OTA91]): 1) computational fluid dynamics for the design of hypersonic aerospace or efficient automobile bodies, 2) computer-based weather and climate forecasts and understanding global environmental changes, 3) electronic structure calculations for the design of new materials such as chemical catalysts, immunological agents, and superconductors, 4) plasma dynamics for fusion energy technology and for safe and efficient military technology, 5) calculation to improve the understanding of the fundamental nature of matter, including quantum chromodynamics, and condensed matter theory, 6) machine vision to enable real-time analysis of complex images for the control of mechanical systems

computers, supercomputers, application programs and software tools, specialized databases, experimental apparatus, digital libraries (books, journals, pictures, sound recordings, films, and other types of information media), computer conferencing systems, bulletin boards, and electronic mail [OTA89].

The HPCC Initiative and the NREN have been the subject of discussion of numerous committees and councils, representing the interests of government, industry, and academia. The Congress of the United States has recently passed legislation that further defines this initiative and determines the responsibilities to be assigned to various government agencies.

This chapter first describes past and current Federal sponsorship in the development of a "national" network for the U.S. "research and education" communities. Although relevant to the development of the networking technology, there is no attempt to describe the multitude of proprietary and special interest networking efforts of the various State and Federal agencies, U.S. industries, and other commercial enterprises. The reader is referred to [QUAR86, QUAR89, QUAR90] for a history of such developments. Second, the current Federal interest in developing the NREN is cited along with a description of recent legislation. Some interesting questions, that ultimately require answers, are raised about the scope, management, and operation of the NREN. Finally, some projections are made about the character of the NREN along with the potential impact on information security.

2.2 Early DARPA Sponsorship

As stated in [HURA90], the ideas for a national research network are traceable to studies by the Rand Corporation in the 1960's. By 1969, the first prototype U.S. network was created by Bolt, Beranek and Newman under sponsorship of the U.S. Advanced Research Projects Agency (ARPA), now called the Defense Advanced Research Projects Agency (DARPA). Sharing of computing resources among researchers was the primary objective. This network, called ARPANET, was an experiment in using leased-line communication media to interconnect a collection of host computers and switching computers. Despite heavy military involvement, the resulting ARPANET turned out to be a fairly open network. It provided the test bed for the development of communication protocols to support functionality such as transmission of graphical data, remote login, file transfer, and electronic mail.

Perhaps the most important aspect of the DARPA development was the research it inspired in packet switching and end-to-end communication across multiple networks, leading to the present-day, widely-implemented Transport Control Protocol/Internet Protocol (TCP/IP). Using "store-and-forward" packet-switching technology, there is no dedicated network path for transmitting a message; rather a message is broken into packets, each of which is independently and dynamically routed through a network, and reassembled at the final destination. The IP protocol serves to connect networks within an internet and the TCP protocol provides reliable end-to-end communications between different machines in an internet.

In response to an overload of traffic on the ARPANET, the Department of Defense in 1983 split off the operation of its military traffic into a separate network called MILNET [MARS89]. The two networks collectively formed what was referred to as the "Internet". Since 1983, with TCP/IP

as the primary protocol, there has been an explosive growth in the number of networks that have connected to and become part of the Internet.

2.2.1 National Science Foundation Sponsorship

2.2.2 Supercomputing Initiative

With funding from Congress, the National Science Foundation established in 1984 a program intended to improve the availability and use of high performance computing to the science research communities.

In 1985-86, NSF selected five sites as National Supercomputing Centers with computing facilities remotely accessible by other researchers through a backbone NSFNET. The selected centers were: San Diego Super Computer Center – located by the University of California at San Diego and operated by General Atomics; National Center for Supercomputing Applications – operated by the University of Illinois; Cornell Theory Center – located at Cornell University; Pittsburgh Supercomputing Center – operated jointly by the University of Pittsburgh, Carnegie Mellon University and Westinghouse Electric Corporation; and John von Neumann Supercomputer Center – located in Princeton, New Jersey.

Promoted as providing backbone connectivity to supercomputer centers, the network traffic of NSFNET was soon dominated by use of general services such as electronic mail, remote login, and file transfer.

With the creation of the Federally funded NSFNET in 1985, ARPANET was eventually phased out and replaced by a new Defense Research Internet (DRI) for unclassified military information that would make use of NSFNET. ARPANET and MILNET became the main constituents of a TCP/IP internet DDN (Defense Data Network) – a subset of the Internet operated by the Department of Defense. Other networks in DDN include DISNET (Defense Integrated Secure Network), SCINET (Sensitive Compartmented Information Network) and WINCS (WWMCCS Intercomputer Command and Control System) of the World Wide Military Command and Control System [QUAR90].

2.2.3 Development of a Higher Capacity Backbone

In an attempt to keep pace with the enormous increases in data traffic (200 million packets/month), the initial backbone network was de-commissioned two years later (1988) in favor of a new backbone [WOLF88]. By July 1988, seven new regional university-based research networks were added and the transmission speed of the backbone was increased from 56 Kbits/sec to 1.544 Mbits/sec (T1). The NSFNET backbone interconnected multiple autonomously administered mid-level networks, which in turn connected to autonomously administered networks of universities and research centers. Multiple peer network infrastructures of other Federal agencies are also connected to NSFNET.

In 1989, NSFNET connected three levels of networks [MARS89, WULF89]:

- the cross-continental NSFNET backbone with 13 gateways (supercomputer site or regional network center) –

BARRNET - Bay Area Regional Research Network (Palo Alto, CA),

JVNCNET - John von Neumann Supercomputer Center Network (Princeton, NJ),

MERIT - Merit Corporation (Ann Arbor, MI),

MIDNET - Midwestern States Network (Lincoln, NE),

NCSANET - National Center for Supercomputing Applications Network (Champaign, IL),

NORTHWESTNET - Northwestern States Network (Seattle, WA),

NYSERNET - New York State Education and Research Network (Ithaca, NY),

PSCNET - Pittsburgh Supercomputing Center Network (Pittsburgh, PA),

SDSCNET - San Diego Supercomputer Center Network (San Diego, CA),

SESQINET - Sesquicentennial Network (Houston, TX),

SURANET - Southeastern Universities Research Association Network (College Park, MD),

USAN - National Center for Atmospheric Research Satellite Network (Boulder, CO), and

WESTNET - Southwestern States Network (Salt Lake City, UT);

- regional networks connecting to the backbone; and
- campus and research organizations.

In 1990, NSFNET consisted of more than 1000 state, regional, and institutional networks, including well over 100,000 computers [GOUL90]. In 1991, the NSFNET backbone was upgraded to 16 nodes operating at 45 Mbits/sec (T3). Its planned protocol base consists of the TCP/IP suite of protocols and also the ISO CLNP (connectionless network protocol).

Numerous other government networks have gateway connections (existing or planned) to NSFNET – including the NASA Science Internet (NSINET), the Energy Science Network (ESNET), and others. In general, the Federal agencies have a vested interest in the Internet.

NSFNET is presently managed by the Merit Corporation under a cooperative agreement with NSF. The operation of NSFNET is contracted to a private nonprofit subsidiary, Advanced Networks and Services, formed by the IBM Corporation, the MCI Corporation, and Merit.

Despite its enormous success, a number of challenging problems remain to be solved, including privatization of the network, priority routing, measurement of traffic and billing, and improved security.

2.2.4 Experimental Gigabit Networks

In June 1990, NSF announced a three-year research effort aimed at funding five test bed experimental networks of gigabit speed. Working in collaboration with DARPA, this was considered a first step in developing a wide-area broadband advanced communication capability.

In 1991, the Office of Science and Technology Policy (OSTP) presented its plan for HPCC in support of the 1992 budget proposed by the Executive Branch of the Government [OSTP91], including funding for the NREN. This plan along with recent congressional legislation calls for gigabit speeds by 1996.

2.3 The Internet

The U.S. portion of the Internet is made up of different parts [CERF91b]. There are Federally subsidized components such as NSFNET, NASA Science Internet (NSINET), Energy Sciences NET (ESNET) and DARPA Test Net (DARPNET) that have agreed to interconnect and carry each other's traffic. There are also commercial networks (PSINET, CERFNET, UUNET/ALTERNET) that are linked together via a commercial internet exchange (CIX) and, via some of its members, linked to the NSFNET backbone. Most midlevel networks are linked to NSFNET and/or commercial networks. International connections have been established through government agreements or through business negotiations by the commercial networks. In all, the U.S. portion of the Internet consists of several government or government subsidized backbones or regional networks, a couple dozen regional/mid-level networks. and thousands of private (industry, university and institutional) networks including private for-profit commercial mid-level and wide-area nets (commercial backbones).

As a first step toward ultimate commercialization[3], commercial networks are presently allowed to establish experimental messaging interconnections to the Internet via one of the mid-level networks. Conditions for interconnection by a commercial network include 1) transporting at no cost to Internet senders (recipients) messages to (from) recipients (senders) on the commercial network, 2) prohibition of use of the Internet for traffic between commercial systems that are not for purposes of research and/or education, 3) prohibition of Internet use for advertisements or solicitations except for services and support for scholarly research purposes – costs to be borne by individual or institutional subscription, 4) optionally making accessible to Internet users any public directory services which assist in identifying the users of commercial services, and 5) bearing any costs associated with physical interconnection. Commercial information providers are permitted to distribute services that support Federal Research and Education, on the Federally-sponsored portions of the Internet.

The global Internet community spans the entire U.S. with links to Africa, Canada, Western Europe, Japan, the Middle East, Australia, Central and South America, New Zealand, and others in the Pacific Rim, Eastern Europe, the USSR, etc. Its growth is so rapid that any estimate of its size is soon obsolete. For example, in 1990, it was estimated that the network included 150,000 connected hosts and millions of users [BENA90, CERF90a]. In 1991, as reported in [CHAR91, ANTH91], the network connected three million users on 350,000 host computers on 5,000 networks in 33 countries. In September 1991, according to [CERF91a], there were more than 5,000 networks connected to the Internet, consisting of more than 570,000 hosts (of which 450,000 are in the U.S. and 90,000 are in Europe).

For some years, the U.S. portion (non-Federal) of the Internet has developed under the informal

[3]Commercialization is defined to be the building of the network through use of commercial telecommunications services whenever feasible.

guidance of the Internet Activities Board (IAB) – a small group of communications experts who volunteered their services, and two subsidiary task forces – Internet Engineering Task Force and Internet Research Task Force [CERF90b]. In June 1991, a user group called the Internet Society was announced to be in operation by January 1992. The purpose of the Internet Society is to function as a professional society, to stimulate interest in and growth of the Internet, to educate the public about the use of the Internet, and to facilitate its continued evolution.

The Internet is expected to continue to operate in its present manner – that is, through its various component networks at various universities, State and Federal agencies, in cooperation with industries and commercial enterprises. According to [CERF91a], the IAB has for several years supported the development of multiple protocol support in the Internet. Presently, numerous protocols operate in various parts of the Internet (TCP/IP, OSI, DECNET, Novell Netware IPX, XNS, etc.). The most common, wide-area protocol is still TCP/IP but the adoption of ISO CLNP as a co-standard appears imminent.

2.4 Present Federal Interest and Initiatives

2.4.1 Federal Councils, Committees and Information Offices

The U.S. Government has since 1984 become increasingly aware of the essential role played by information technology in practically all areas of research and development. Bell reports [BELL88], the U.S. Congress requested in 1986 that OSTP study the potential development of a communications network for research computers, including supercomputers at universities and Federal research facilities. In recent years, a number of agencies, committees, and councils have been instrumental in the advising the government in its planning for HPCC and national networking –

1. Office of Technology Assessment (OTA) - advises the U.S. Congress on matters concerning science and technology;

2. Office of Science and Technology Policy (OSTP) - advises the President branch on matters of science and technology and coordinates interagency issues regarding science and technology; advice regarding HPCC and NREN (as a component of HPCC) is formulated primarily through its Federal Coordinating Council on Science and Technology (FCCSET); it will be assigned major responsibility in developing national plans for HPCC;

3. National Science Foundation (NSF) - a U.S. government funding agency charged with advancing research in science and technology;

4. President's Council of Advisors on Science and Technology (PCAST) - advises the President on matters of science and technology (membership is from the private sector);

5. Federal Research Internet Coordinating Committee (FRICC) - superseded by FNC (see below), this informal committee was formed to coordinate U.S. Government support for the development and use of the Internet; government agencies initially included the Department of Energy (DOE), the Defense Advanced Research Projects Agency (DARPA), the National Aeronautics and Space Administration (NASA), and the National Science Foundation (NSF), along with observers from the Internet Activities Board; and

6. Federal Networking Council (FNC) - the successor to and enlargement of FRICC, this is an independent interagency group; it coordinates the use of the U.S. portion of the Internet by Federal agencies and provides liaison with OSTP; some members of the FNC are also members of the Coordinating Committee for Intercontinental Research Networks (CCIRN); the FNC has representatives from numerous Federal agencies – OMB (Office of Management and Budget), NSA (National Security Agency), DISA, NOAA, DOE, DARPA, HHS, OSTP, NIST, EPA (Environmental Protection Agency), USGS, GSA (General Services Administration), NTIA (National Telecommunications and Information Administration), NASA, NSF and the Department of Education; advisory committee members come from the IAB, higher education, national research laboratories, computer and communications corporations, and private industry.

2.4.2 High Performance Computing and Communications Initiative

Recent reports issued by OSTP and OTA [OSTP91, OTA91] address the HPCC requirements to sustain and extend U.S. leadership in all advanced areas of computing and networking and to meet the so-called "Grand Challenges".[4] Congressional legislation [CONG91b] identifies four components of the program:

1. hardware - development of more powerful supercomputers and networking technologies;
2. software - development of higher performance software to effectively apply the power of supercomputing;
3. education and basic research - training of computational scientists to effectively use supercomputing technology and training of computer scientists and engineers to develop new supercomputer hardware and software; and
4. networking - deployment of a national computer network capable of transmitting information at multi-gigabit speeds to allow for the appropriate communication and access to shared resources.

2.4.3 Mandate for a National Research and Educational Network

Both the OSTP and OTA identify the development of the NREN as an essential component of a HPCC program. Such a network is required to provide distributed computing capability to the U.S. research and educational community (government agencies, industry and research laboratories, universities, libraries) and would further advance research on very high-speed networks and computer applications. In this role, the NREN component will dramatically expand and enhance the research and educational aspects of the U.S. portion of the larger Internet. The intent is to provide a wide-spread, uniform, high-performance (gigabits/second) national infrastructure and also provide

[4] A layman's list would include 1) forecasting severe weather events, 2) human genome research, 3) predicting new superconductors, 4) air pollution, 5) aerospace vehicle design, 6) energy conservation and turbulent combustion, 7) microsystems design and packaging, 8) predicting directions and consequences of changes in the earth's biosphere, 9) development of gigabit networks, and 10) education using a national network.