NSF Implementation Plan for Interagency Interim NREN

Robert Aiken (NSF)
Hans-Werner Braun (SDSC)
Peter Ford (LANL)
Editor: Kimberly Claffy (SDSC)

1 May, 1992
printed: May 15, 1992

Abstract

This document outlines an architecture and implementation plan for the National Science Foundation's Interagency Interim National Research and Education Network (NREN) component of the HPCC Program. The Interagency Interim NREN is intended as a near term research and development program to extend the capabilities and breadth of connectivity of today's research and education (R&E) networks. This network will provide not only the opportunity for agency backbones and mid-level networks to connect to NSF-funded networks, but also act as a limited testbed for new technologies in the migration to a gigabit NREN. In order to ensure a smooth transition to a national gigabit network infrastructure which can support the combined requirements of federal agencies, the U.S. research and educational institutions, and industrial collaborations with U.S. R&E institutions, the National Science Foundation has collaborated with other federal agencies in establishing this framework for the Interagency Interim NREN.

Contents

1 Introduction

The National Research and Education Network (NREN) is a component of the Federal High Performance Computing and Communications (HPCC) Program as defined by the High Performance Computing Act, 1991 (PL 102-194) and the High Performance Computer and Communications Report of the Committee on Physical, Mathematics, and Engineering Sciences Federal Coordinating Council for Science Engineering and Technology. The intent of the NREN is to interconnect the nation's education and research communities at gigabit per second data rates, thereby facilitating research and development collaborations and enhanced access to information resources and computational capabilities. The relationship of NREN to the goals of the HPCC are detailed in "Grand Challenges 1992: High Performance Computing and Communications, the FY 1992 U.S. Research and Development Program" and the equivalent report for the Fiscal Year 1993 budget (respectively, The Blue and Teal Books). The development of NREN will be accomplished by two supporting NREN activities:

- The Interagency Interim NREN will result from the evolution of the current networks which interconnect and serve the U.S. research and education communities. This infrastructure includes federal networks such as the NASA Sciences Internet, the DOE Energy Sciences network, the NSFNET backbone and the DARPA advanced research networks, in addition to the mid-level networks which are interconnected by the NSFNET backbone. The mid-level networks include the regional networks developed under NSFNET sponsorship or guidance and "community of interest" networks, such as HEPnet, developed under federal agency sponsorship to connect sites in the research and education community. A key activity of the NREN program is to enhance the interconnection technologies and strategies for both federal and non-federal networks, without interfering with the autonomous management of each component network. The ultimate goal of NREN activities is an effective and richly connected Internet, appearing logically as a single network serving the needs of the research and education community, much as the matrix of U.S. telephone companies appears as a single transparent system to most users.
- The goal of the Gigabit Research and Development program is to spur the development of communication technologies used to build wide scale gigabit per second networks. This program will foster collaborations among industrial, academic and national research and development efforts to develop, deploy and evaluate these new technologies in network testbeds. As these technologies mature they will be integrated into the Interagency Interim NREN.

The National Science Foundation is responsible for a broad set of activities and tasks in the Interagency Interim NREN program. Federal agency R&E networks are strictly responsible for end-to-end connectivity among research staff and sites directly funded to accomplish agency programmatic goals. The NSF network program complements these efforts by providing general infrastructural network support in addition to providing connectivity for NSF sponsored R&E projects. Since many of these "gap filling" functions are essential to the reliable provision of connectivity to meet agency programmatic goals, NSF's role as a default network requires coordination with the other federal agency network programs. NSF's Interagency Interim NREN activities include:

- The NSFNET program must address the network connectivity requirements of the general research and education community. Responsibilities include supporting the end-to-end connectivity requirements of all R&E researchers, including those of federally funded research facilities as well as those of staff working on projects not funded directly by the NSF.

- The NSFNET program funds connectivity requirements at several levels including campus to regional networks and inter-regional connectivity. Over time, NSF will target its funding to those campuses which have financial impediments to connecting into the U.S. Internet. In a similar manner, NSF will target funding of inter-regional connectivity. NSF will initiate programs to fund new connectivity requirements including very high bandwidth requirements, innovative community of interest programs including projects addressing access to Libraries and Primary/Secondary Education , and additional requirements which fall under NSF's role as "gap-filler". The NSF expects that over time network subscribers will increasingly bear the costs of their network connections, as the availability of these network connections becomes available in the commodity telecommunications market.

- The Interagency Interim NREN will continue to rely on the mid-level networks for satisfying end user and site connectivity requirements. The NSF will increase the capabilities of these networks where required, and in addition encourage the development of tools and plans to provide for robust, reliable operation of the autonomous R&E networks in a coherent yet distributed manner.

- The Interagency Interim NREN will be an open system based on standardized protocols and accepted network operating principles, and will interconnect a wide variety of networks including federal backbones, mid-level networks, community of interest networks, and international networks supporting global research and education communities. Commercial network providers which carry traffic and provide network access in support of the U.S. R&E community will also interconnect to the Interagency Interim NREN. NSF will support the development of a standard interconnection architecture in support of goals such as, but not limited to, the development of Network Access Points (NAPs), where networks requiring access to U.S. R&E networks can connect. The NAPs, where the NSF will provide for routing coordination and management, will be the locus for the evolution of the NREN interconnection architecture. The NAPs will allow for interconnection of federal, mid-level, international and commercial networks and will be implemented in a manner which is not restricted by the NSF Appropriate Use Policy. The NAPs will initially be funded and managed by the NSF; it is a goal for the NAPs to become funded and managed by the participating networks over time.

- The rate of growth of autonomously operating networks supporting the research and education community is beginning to strain the integrity of the current Internet. The NSF is committed to the development of technologies and deployment strategies to overcome the scaling limits inherent in the IP framework on which the current Internet architecture is based. In this context, the NSF is also committed to the NREN goal of a robust multiprotocol network, which in addition to using IP technologies, will use OSI datagram technologies (CLNP, IS-IS, IDRP, etc.).

- Merit, in line with a cooperative agreement with NSF for the operation of the NSFNET backbone, maintains a global policy routing database for the Internet. NSF will continue to support this functionality and will actively pursue methods for distributing this responsibility.

- The NREN aims to support R&E collaboration and network access across the nation. Success will require secure operation of the network, implying a need for authentication of users, sites, and computing and routing resources. These authentication mechanisms will allow autonomously operating networks, sites, and users to enforce authorization decisions appropriate to their specific operating environments. Sites and networks may also choose to use these authentication mechanisms in support of accounting for network and computer resource utilization. The NSF will facilitate the deployment of standardized security architectures and infrastructures for the NREN in conjunction with federal efforts in the security arena.

- The NSF will support the development and operation of network information services including the infrastructure elements required for assigning network resources such as network numbers, humanly readable names, points of contact and routing table maintenance. The infrastructure will be distributed in nature, taking advantage of standardized services and protocols such as DNS, X.500, Z39.50 and hopefully spurring innovation in distributed information management and access systems.

- The NSF will support the advanced research and development of software tools and systems to fully utilize the rich connectivity provided for by

NREN, in concert with other programmatic efforts in the HPCC program. Relevant activities include access to remote computational centers, digital libraries, remote educational resources and R&E staffs across the nation.

All of these activities will facilitate access by the R&E community into digital libraries and data repositories; widespread access to large scale distributed computing resources; interrogation, retrieval, and visualization of data from data bases; remote control of experiments at unique national facilities; and multi-media, multi-site teleconferencing to foster increased collaboration and education opportunities. The NREN implementation of this infrastructure must satisfy three general requirements:

- In order to ensure a smooth transition to a national gigabit network infrastructure which can support the combined requirements of federal agencies and research and educational institutions, the National Science Foundation (NSF) must collaborate with other federal agencies in establishing a framework for the NREN.
- The initial deployment of this network, the Interim Interagency NREN, must allow agency backbones and mid-level networks to connect to NSF-funded networks. Commercial and International networks will also connect when meeting requirements of the U.S. R&E community.
- The major federal agency networks: ESnet, DARTnet, TWBnet, NSI, and the NSFNET backbone, should upgrade as necessary in line with NREN goals, while maintaining a stable operational environment and satisfying their own HPCC and other R&E requirements.

The goals and requirements of the NREN are considered to be a proper subset of, and therefore consistent with, the overall NSFNET programmatic goals.

This document provides background information on these relationships and issues. After briefly describing the current architecture of the federally funded R&E networks, it outlines the requirements for advancing the NREN infrastructure, and milestones for effecting this plan.

2 Current Architecture

The Interim Interagency NREN will evolve from the federally sponsored U.S. portion of the Internet. The NREN program will leverage the strengths of the existing national infrastructure consisting of NSFnet, NSI, ESnet, TWBnet and the Federal Internetwork eXchanges (FIXes), as well as the entire Internet hierarchy of mid-level, regional, and campus networks. Each of these networks operates autonomously, serving its own community of interest, but agrees to

Figure 1: Network Hierarchy

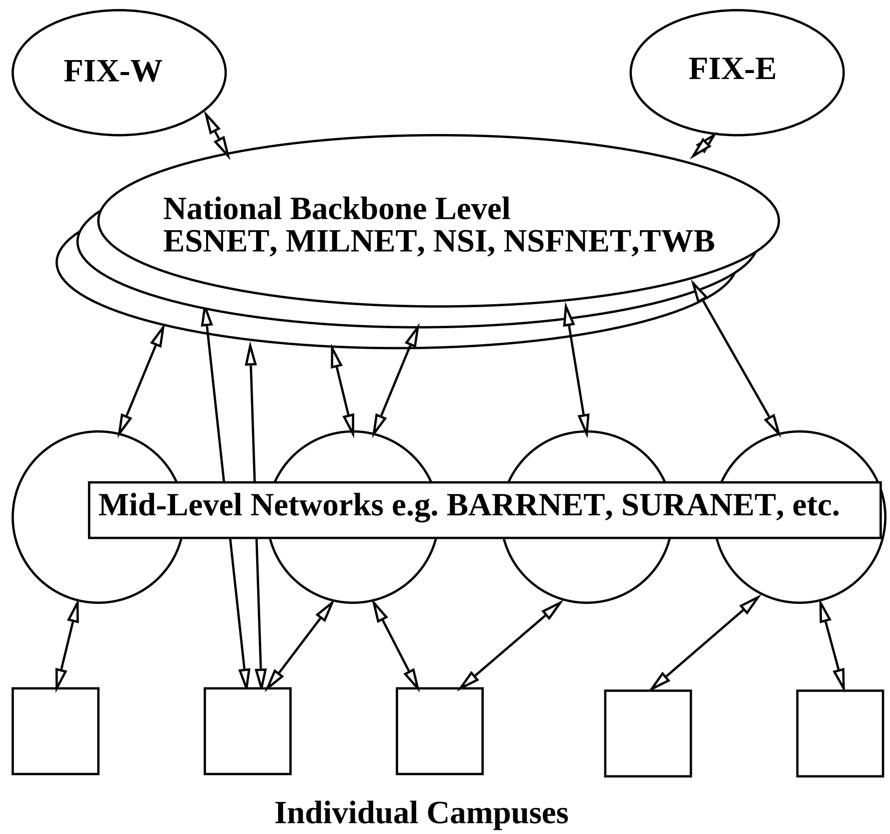

interconnect and exchange traffic to form the U.S. Internet and the operational base for the Interagency Interim NREN.

The U.S. Internet is loosely based on a three level hierarchical model that is national in scope. Agency backbones, mid-level networks, and connected local site networks are the levels in the hierarchy. The top level consists of the agency backbones, which currently interconnect at the Federal Internet Exchanges (FIXes) at NASA Ames and the University of Maryland. (See Figure 2.)

This abstraction is not wholly accurate, as it ignores commercial network providers, international networks, and interconnections which bypass the strict hierarchy. For the sake of this discussion, commercial Internetwork providers and international networks best fit into the mid-level network abstraction, since they use the NSFNET and agency backbones to connect their customers to the U.S. R&E community. Agency backbone networks other than the NSFNET backbone may and do directly connect to mid-level networks and campuses to meet their mission critical requirements.

There are three fundamental elements to building the U.S. Internet hierarchy: routing information, actual switching of network packets (datagrams) between networks, and transit of packets within and between networks. NSFNET has a very special role in the hierarchy: it acts as a generic transit, routing, and switching network for R&E networks. NSFNET maintains a global routing database and accepts packets for transit to any destination R&E network. This significantly eases the operation of mid-level networks since they can count on pushing traffic up into the NSFNET for transit to other mid-level networks. Agency backbones also use the default service of the NSFNET to obtain connectivity to networks and sites to which they are not directly connected.

2.1 Agency backbone networks

2.1.1 National Science Foundation: NSFNET

Merit Computer Network operates the National Science Foundation's backbone network in a cooperative agreement with Advanced Network and Services (ANS), MCI, and IBM. The NSF backbone interconnects 16 mid-level networks and the NSF supercomputer centers at 45 Mbits/sec. The cooperative agreement has been extended for up to 18 months past November 1992. The NSF is preparing to recompete the provision of R&E "backbone" transit services; the awards are planned to be made prior to April 1993.

2.1.2 Department of Energy (DOE)

The Energy Sciences Net (ESnet) activities, which provide for network connectivity to major DOE Energy Research users and sites, are part of the overall NREN effort. ESnet supports DECNET and IP services and is in the process of including support for OSI services. DOE activities in high performance SMDS packet switching and cell- switching technologies, as well as other NREN activities, are part of ESNet evolution. ESnet is managed and administered by the Lawrence Livermore National Laboratory's National Energy Research Supercomputer Center (NERSC).

2.1.3 NASA Sciences Internet and NASA AEROnet

The NASA Science Internet (NSI) encompasses TCP-IP and DECNET activities that satisfy programmatic requirements in NASA. NASA is planning to collaborate with the DOE on a national pilot project based on SMDS technology. AEROnet supports the programmatic goals in computational aerosciences of NASA's NAS project at NASA Ames Research Center and interconnects NASA centers and commercial aerospace and aeronautical firms. AEROnet and NSI interconnect and gateway internetwork traffic at NASA Ames.

2.1.4 DARPA networks

DARPA supports DARTnet and the Terrestial Wideband Network (TWBnet), research and development networks for the development of new protocol, switch and router technologies. Research and development on these testbeds aims to support a wide range of advanced network services such as packet video, multicasting, distributed multicomputer simulation, network resource management and control.

2.1.5 International Networks

The NSF supports extensive network connectivity to international R&E sites and networks. This is accomplished through direct funding of international connections as well as through interagency cooperatively funded links.

2.1.6 Federal Agency networks

The National Institutes of Health, National Library of Medicine, Environmental Protection Agency and Department of Education networks and network based projects will coordinate and interconnect with the Interagency Interim NREN. Federal agencies beyond the ones listed above will also be accommodated as required. The NSF is responsible for making this interconnectivity possible.

2.2 Mid-level Networks

The logical layer below the NSFNET and other agency backbones includes the mid-level networks.[1] Mid-level clients connect to the NSFNET backbone and may also connect to agency backbones and to other mid-level networks. Mid-level networks are responsible for connecting sites such as academic, educational and corporate institutions. Most mid-level networks have evolved from networks which originally connected geographically proximate university sites, or university sites affiliated with NSF supercomputer centers. Over the last years, mission agencies such as DOE and NASA have also relied on mid-level networks for the majority of their connectivity to university-based researchers, either using direct connectivity to individual mid-level networks or indirectly connecting at the FIXes via the NSFNET backbone.

[1] Mid-level networks have also been called *regionals* , a term which at one time reflected their geographical span, but we will use the term *mid-level* to reflect its hierarchical position in the architecture.

2.3 Campus Networks

The most important end components of the network are the individual sites which include university and college campuses, research laboratories, private companies, educational sites such as K-12 school districts, etc. The aggregated investment at these network distribution sites dramatically surpasses that of federal government investment in backbone and mid-level networks. For example, the annual cost of operation of an individual large campus network can approach that of an entire government agency T-1 based backbone. These site networks form an integral part of the overall infrastructure, largely financing their expenses out of internal funding, and are quite independent in both internal as well as interconnection decisions.

2.4 Commercial Networks

Collaboration among U.S. industry, federal research staffs and academic research is critical to the success of the HPCC program. The Interagency Interim NREN will facilitate network connectivity to commercial sites and networks where appropriate. The NSFNET backbone will aid in this effort by providing interconnection sites, routing information, and switching and transmission capability so that commercial networks can gain access to U.S. R&E networks.

3 Advancing the Infrastructure

The NREN community is facing serious scaling problems due to the success and the resultant massive growth of the Internet([Lottor 92]). The Internet used to be managed and run by a small group of U.S. research organizations who agreed to simple rules of interconnection and had plenty of resources, IP address space and routing table entries, to use. The current Internet consists of more than 5000 networks spanning the globe, operated and managed autonomously, with most of these networks using the U.S. network infrastructure as the focal point for global routing and intercontinental transit. The NSFNET backbone providers (Merit and their collaborators) are currently responsible for routing coordination and management for the entire global Internet. This growth is placing tremendous pressure on current routing and router technology since it was never anticipated that it would be necessary to manage as many as 5000 networks.

There are also scaling problems due to the availability of new technologies. In the past the bandwidth for most wide area links was comparable, but now these links commonly show great diversity ranging from 10-60 Kb/sec to 50 Mb/sec data rates. It is likely that gigabits/sec network links will be available within the next 2 to 3 years. Switched network services will allow the network user to dynamically add links on demand or to modify the bandwidth available. It will

Figure 2: Network Service Layering

Services above the network layer(IP), such as Network Information, Electronic Mail, Authentication Services, Nationwide File Systems, etc. Connection services such as TCP/TP-4.	layer 3++
Internetwork Packet connectivity (IP, CLNP, etc.). Internetwork Routing, and Internetwork operations.	layer 3
Packet transport, switching. Public data services such as SMDS, ATM, Frame Relay, X.25. Links interconnected with switching nodes.	layer 2
Bit transport such as DS-0,1,2,3 (T-carrier) and Sonet, SDH.	layer 1

be a challenge for the NREN to accommodate the growth in complexity due to the diversity of network bandwidths and richer interconnectivity of networks.

The Internet used to serve only the research community, but has recently evolved to service a wide variety of network connectity including commercial, research, education and governmental users and sites. Thus, the distributed management of the network must evolve to accommodate the requirements of new Internet users and providers at the same time meeting the ever changing objectives of the R&E community.

Figure 3 illustrates the layering terminology used in subsequent discussions of network infrastructure. It is loosely related to the layers used in discussing OSI network architectures.

Although currently an IP-based packet switched network, the NREN must gradually assume increased support and use of OSI protocols. Media and transmission independence is a key feature of the current Internet architecture, which accommodates a variety of technologies spanning wide ranges of bandwidth, latency, reliability and availability. Such flexibility on the part of the existing infrastructure indicates that the Internet can indeed evolve to include new technologies and service offerings of the telecommunications industry.

3.1 Transmission Requirements

3.1.1 Multiple Scaling Requirements

Historically, NSFNET has provisioned network bandwidth for the aggregation of many users into appropriately sized communication pipes. Recently, the national infrastructure has proven to be an enabling technology for applications which are highly demanding of network resources. Some of these applications such as scientific visualization require predictable high bandwidth to succeed. Maintaining an environment which can accommodate these highly demanding real-time applications will be challenging as the network continues to scale in bandwidth and ubiquity, resulting in rather complex metrics for network performance. The Interim Interagency NREN will scale along multiple dimensions while securing predictable high bandwidth paths to individual working environments.

We are clearly in the decade of multimedia computing and information access. The current Internet rarely focuses on isochronous services such as audio and video streams, and the current connectionless packet transport does not offer the capability of reserving and managing guaranteed bandwidth for such applications. The NREN will face formidable challenges in accommodating such services. SMDS and B-ISDN, the packet services which most telecommunications companies plan to offer in the mid-90s, hold promise is being able to handle isochronous traffic. Supporting multimedia technologies via hard multiplexing of channels in the network does not effectively share bandwidth and is not scalable. The NREN program should strive for switching architectures and routing support to accommodate digital traffic generated by continuous media applications such as scientific visualization, remote steering of simulations, retrieval from digital video archives, remote learning, and video conferencing.

3.2 Scalability and Administered Control

The last several years have brought a dramatic growth in network numbers in the Internet, a trend which continues to accelerate in domestic as well as international arenas. (See Figure 3.2.)

Embracing dramatic developments of new technologies, while maintaining reliable ubiquitous service to the wide range of ever-increasing number of customers, poses a serious challenge to the evolution of the network infrastructure. The NSF is currently working with a small exploratory group of internetwork engineers and scientists, the ROAD group, to evaluate strategies for scaling the Internet routing and addressing technology base in light of the growth in the number or networks and end systems.

In the long term, fundamental limits in the size of the IP address space must be overcome. One approach is to transition the Internet to use CLNP (ISO-IP) as the network packet datagram. The immediate Internet scaling problems do not stem from fundamental IP limitations, but rather from the use of a single

Figure 3: Growth in number of IP Networks

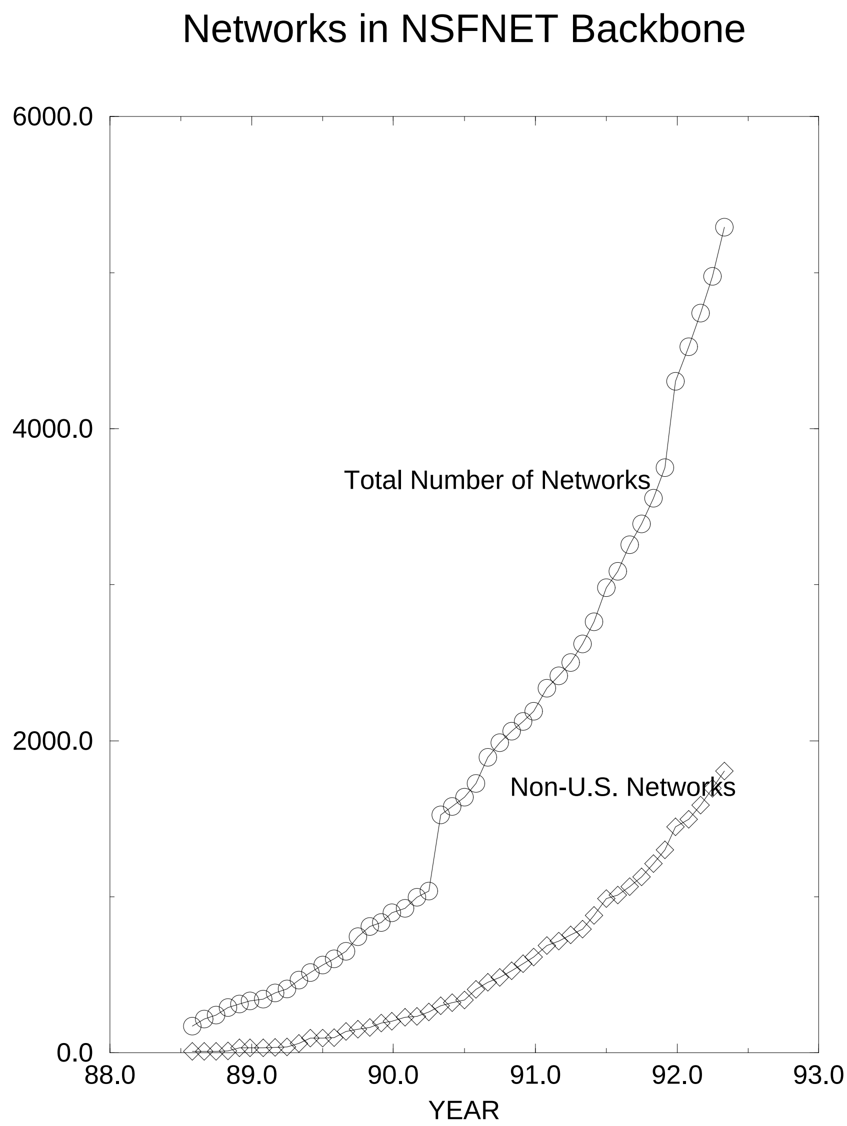

level hierarchy (class A, B, C) to assign network numbers, and the use of this flat network number space for routing. The multilevel hierarchical addressing scheme for OSI NSAP addresses scales more effectively, at the expense of having much larger network addresses.

The future assignment of IP networks addresses needs to address the needs of networks which are larger than the 255 end systems managed under a single class C network, but are significantly smaller than the 65,000 end systems handled by a class B network. It is also the case that the current routing system does not have an effective way to compress the routing data base along topological lines, since network addresses are not assigned in a manner which can encode routing topology information. These scaling problems need to be addressed by IP interdomain routing and an effective plan for future assignment of IP network numbers which takes the routing topology into consideration. The current network number address space is not hierarchical, being logically flat (ignoring subnetting since interdomain routing is based on network numbers). Thus, major transit networks, like the NSFNET backbone, are required to maintain a complete enumeration of each individual network's route in the Internet. This is memory intensive and also requires significant routing information exchanges between networks. The ROAD group and the IETF are working on future routing and addressing plans for the Internet.

3.3 Chaos Into Order: Network Access Points

The logical topology of the interconnected federal agency backbones and mid-level networks is the direct result of the exchange and propagation of routing information in a well coordinated manner using standard Internet routing protocols and a moderate amount of route filtering. To date, there has been a strong coupling between administrative policy and internetwork topology, as seen in the four major federal agency backbones (NSFNET, TWBNET, ESNet, NSI, and the Milnet) and the Federal Internet eXchanges (FIXes). In general, networks interconnect with each other to exchange routing information and to interchange traffic and as a result the networks are called peer networks and the "peer" with each other. The routing information a network (A) sends to a peer network (B) indicates networks and sites that B can reach by sending traffic through network A. A network might selectively advertise its own available routes to a peer network since it may not want to act as a transit for all the routes it knows; this frequently occurs when a network has "private" connections to other networks which are used only to meet certain connectivity requirements. Similarly, a network may choose to selectively take in routes advertised by a peer network choosing instead to use routes discovered from other peers, or simply not to route to those specific destinations. This route filtering is often done to enforce policies such as transit restrictions, or to indicate route preferencing such as high vrs. low speed connection preferences. The Agency backbone networks currently peer with each other at the FIXes. Although not exactly a separate architectural layer, the FIX peering points overlay the existing three-layer architecture of the U.S. Internet.

Conceptually, FIXes facilitate robust connectivity and the implementation of well-defined interagency routing policies. Route filtering is done between the backbones to control the routing and packet forwarding for sites and regional networks connected to multiple backbones. Unfortunately, this filtering also inhibits sharing available bandwidth and eliminates the ability to use redundant connectivity to repair network partitions which may occur due to backbone network failures. Routing technologies which are more dynamic, such as the Border Gateway Protocol[Lougheed 91] which is in the early stages of deployment, can facilitate flexible management of routing information in place of strict route filtering.

Agency networks (e.g., NSI, TWBnet, ESnet) define agency-related traffic as their objective in contrast to the NSFNET backbone, and its associated mid-level networks, that assume responsibility for transmission of any R&E traffic. Although sites and campuses may connect to multiple transit backbone networks, the federally established NSFNET plays a special role as a default transit network. Merit as part of its responsibility in providing the current NSFNET backbone also acts as a "routing arbiter" and has historically provided all networks with a picture of the routing topology of the Internet. This routing arbiter function has become significantly more difficult over time due to the explosive growth in the number of networks, and to the evolution of the Internet supporting more than simple R&E network traffic. The NSFNET accepts transit traffic which complies with the NSFNET Appropriate Use Policy (see Appendix A) and it is Merit's responsibility to enforce the NSF AUP.

However, the model of NSFNET as a generic transit provider for the Internet is breaking down. Commercial network providers now offer inter-regional connectivity for traffic which does not comply to the NSF AUP(see appendix A). In fact, the drive for commercial network interconnectivity has led to the formation of commercial interexchange points (CIXes), analogous to the FIX concept, with national scale commercial backbones like Alternet, PSInet, and Sprint, as well as commercial mid-level networks like CERFnet using the CIXes for exchange of non-AUP compliant network traffic. Unfortunately, this has exposed several weaknesses in the Internet routing technology base, which was not designed with the NSF AUP in mind.

As private vendors assume an increasing portion of the now federally-subsidized market for wide area internetworking services and infrastructure, backbone clients will eventually contract with commercial transit networks for inter-regional connectivity. At the same time, mid-level networks are gradually seeking greater self-sufficiency, and NSFNET itself, in its role as the general purpose transit network for mid-level networks and networks of other agencies, plans to seek connectivity from more than one network service provider. These transitions in the structure of the Internet will allow multiple providers to compete for both mid-level and commercial network customers, increase the number of viable alternatives for network clients, and foster competition in the entire industry. This should result in a competitive cost structure which plays into the economies of scale available in the telecommunications industry. The net benefit to the R&E community should be the broad availability of low cost, high bandwidth, professionally managed Internet connectivity.

Distributing the transit function across multiple networks poses implementation and management difficulties unless fundamental routing problems are solved. Current routing technologies work best in a pure tree hierarchy (see [Rekhter 89]), which is not an appropriate topological model given the scope, scale and interconnectivity of current national network providers. Indeed, the advent of the Border Gateway Protocol[Lougheed 91] is due to members of the NSFNET project, in collaboration with similarly interested Internet engineers, perceiving the need for the Internet interdomain routing architecture to evolve away from a strict hierarchical routing topology. Wide scale deployment of BGP will facilitate migrating the Internet routing topology over the next few years. The evolution of the current logical interconnection architecture from an abstract tree toward a mesh of multiple connections among elements requiring differing levels of service quality and sensitivity may lead to asymmetric routing and transit paths. This is not a significant problem, providing that the primary service providers are of similar quality in an operational sense and provide comparable bandwidths, which seems reasonable for the providers of NSFNET backbone connectivity. A key to moving from a strict hierarchy to a mesh connected routing and network traffic flow topology is to build an interconnection architecture which facilitates the transition. It is unreasonable to expect networks to convert routing technologies overnight and on a single "flag day". By implementing a network interconnection plan which can accommodate both the old and new architectures, and is able to "map" between the two, the transition, which NEG expects will take at least 2 years, can be accomplished. This interconnection plan will be funded and managed by the NSF with an eye for transitioning to a future where the participating networks fund and manage the NAPs.

Any viable interconnection strategy must allow Internet network providers, including non-federal networks, reasonable access to Interagency Interim NREN. The responsibility for building a coherent picture of the routing within the Internet, which currently falls under the auspices of the closed federal R&E community at the FIXes, should diversify towards a system where all network providers participate in shaping the routing fabric at Network Access Points (NAPs). In this section we briefly discuss the functionality of a NAP, please see [HWB 92] for a more complete discussion and set of definitions.

A NAP must accommodate a wide range of service providers, both from the R&E and commercial communities. [2] NAP's will allow federal networks to

[2] FIXes can be viewed as a special case of NAPs in the current environment, which focus on federal networks.

Figure 4: Network Access Point

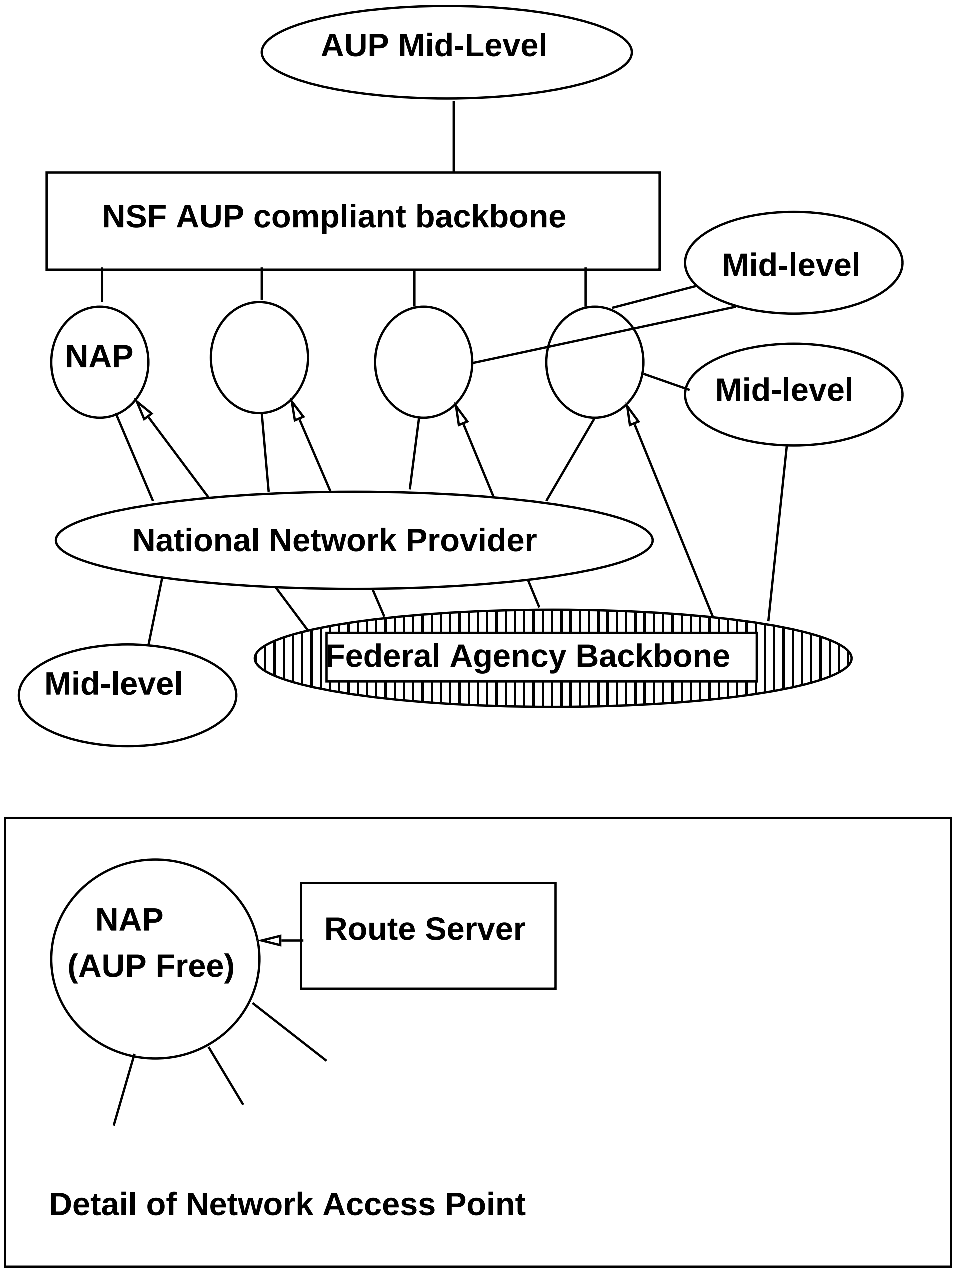

peer with each other, as well as with non-federal and mid-level service providers, in support of the R&E community. NAP's will also allow mid-levels and non-federal service providers to peer with one another, which will enhance and extend the national infrastructure and connectivity without affecting the federal networks. Properly designed, implemented, and coordinated NAPs would provide an evolutionary path for the NREN, NSFNET, and Internet, in line with the goals of the HPCC. Initially, NSF will manage the NAPs through a collaborative agreement for a "routing arbiter". NAPS, at a minimum, must:

- Integrate and interconnect different network technologies, services, protocols and routing strategies.
- Provide a mechanism for new network service providers, including commercial networks, to connect to and peer with the R&E oriented infrastructure.
- Improve the robustness of the system in the face of the increasing complexity of interconnectivity requirements.
- Allow for the evolution of network layer routing in the Internet.
- Monitor and ensure correct routing exchanges between peering networks.
- Provide a mechanism for generating default routing information for networks which do not want to carry the entire routing topology of the Internet withing their network.

The NAP attachment points, in conjunction with network customers, should develop a minimal set of standard routing protocols and a standardized methodology for sound routing and transit of packets in the NREN. All networks connected to the peering points would subscribe to a base set of operational requirements. A NAP based system does not preclude networks peering at a non-NAP site and exchanging peer-wise specific routing and traffic.

The NAPs must be implemented in a robust and reliable manner. Robust operation will depend on the existence of multiple and geographically dispersed NAPs, critical networks interconnecting to multiple NAPs, sound operational support and stability of the physical environment, and a robust architecture for routing between the various administrative domains.

3.3.1 Other Benefits of Route Peering

Peering provides for critical NREN capabilities: a mechanism to interconnect networks of varying speeds (e.g. 45Mbps, 1.5Mbps, 56kbs), technologies (e.g., SMDS and ATM), and dissimilar policies (e.g., DOE, NASA, NSF, PSI, ANS). Furthermore, the concept of controlled interconnections allows the phased introduction of new technologies, including cell-switched ATM, SMDS and SONET style networks, into an operational gigabit NREN.[3] Multi-protocol routers may peer with each other at the interconnection sites on a protocol specific basis, allowing the incremental introduction of OSI and other new protocols.

Although route peering generally occurs at layer 3 as discussed above, the need for controlled interconnections will extend to other layers with the introduction of new services. Peering at the application level, for e-mail or directory services, will be of interest as well. Traditionally, such peering occurs through application gateways, with often no more than a bilateral agreement to cross-reference and connect, via pointers, aliases, and disparate name spaces and management domains. Operationally feasible multilayer interconnections will require careful design and a well-understood architecture.

3.4 Enhancements to the Infrastructure

3.4.1 Security

Both the Congress, Public Law 102-194, and the Executive Branch, "Grand Challenges: High Performance Computing and Communications" reports for FY 92 and FY 93, have stated the need to provide enhanced security within the Interagency Interim NREN. Secure identification of principals on the network (e.g. users, computers, routers, switches, etc.) is a requirement. Once the principals are authenticated, it is possible to provide for different levels of security and end-to-end reliability as required on an application by application basis.

There are several candidate authentication technologies which might address this requirement. Authentication is also required as a basis for future accounting mechanisms which will be used for capacity planning and billing purposes. Possible authentication mechanisms include, but are not limited to: the El Gemal algorithm advocated by the NIST for Government use, "Smart card" technologies, and RSA's public key algorithm identified by many OSI and IP applications (e.g. X.500 and privacy enhanced mail). There is a need to evaluate these authentication schemes in light of Interagency Interim NREN requirements and develop a deployment strategy.

Initially, NREN requirements for enhancing security should focus on securing NREN routers and switches and their access by authorized network managers and network mangement tools. In addition, resources which might have limited access requirements, such as supercomputer centers and limited access databases, will use these authentication tools to make authorized access decisions. Progress in these areas will facilitate securing applications such as electronic mail, file transfer, remote logins and other distributed applications.

Specific areas of investigation include:

- Evaluate authentication methods appropriate for use in the NREN, including certificate distribution and revocation

[3] Please see 5 for reference to the evolving DOE/NASA project to create a national SMDS pilot testbed network.

- secure supercomputer and research center access
- predictable and controllable routing
- secure network management and analysis tools (e.g. SNMP).
- secure e-mail, both SMTP and X.400 technology bases.
- secure directory services (e.g. DNS and X.500 enhancements).
- secure remote login services
- a distributed and delegated system for security center and Emergency Response interaction that would include the Federal research and education networks, CERT , CIAC, mid-levels, commercial service providers, and campuses.
- encouraging vendor development and deployment of secure hardware and software packages, including "user friendly" security and risk analysis software packages which can educate both end system managers and users on how to secure their systems.

3.4.2 Network Information Infrastructure

The current Internet provides inadequate information about itself and the sites and users connected to the network. Unlike the telephone system, there is not a commonly accepted and understood way to find out simple information, such as an electronic mail address of a Internet user. As one of the largest distributed systems in the world, the Internet has remarkably little in the scope of distributed query and information access tools, such as a directory service, to navigate the enterprise. It is easy to imagine that Internet users might want to have more than simply their name and e-mail address in a directory service, and they might want to choose to advertise their occupation, professional affiliations, areas of interest, favourite fine lunch or dinner site, etc.

To date, the principal distributed directory service deployed in the Internet is the Domain Name System (DNS). The DNS is primarily responsible for providing host name to host network address mapping (and its logically reverse mapping). The DNS also maps domain names to site locations where mail should be delivered (the Mail eXchanger(MX), function). The distributed nature and deployment of the DNS is crucial to the growth of the Internet. Centralized management and control of this information is not possible given the autonomy of administration of networks within the Internet. The mail exchanger function allows the concept of the core service of the Internet, e-mail based on SMTP, to be extended to sites that are not directly connected to the Internet which has proven to be extremely valuable. The functionality provided by the DNS is critical. The current quality of DNS implementations, mostly derived from the Berkeley BIND software base, is improving, but older versions of

BIND generate a large amount of spurious traffic on the NSFNET[Danzig 92]. Vendors of software need to be encouraged to deploy current versions of the software.

There are several enhancements to network based information services needed to facilitate the usage of the Interagency Interim NREN. It is assumed that a large effort will be made in the near future to implement widely available "white and yellow pages" services which are based on deployment of "The Directory" ,X.500[ISO 9594-1], which is the OSI technology meant for building global scale directory services. Functionally, the Directory is accessed by Directory User Agents (DUAs) and the Directory Information Base (DIB) is maintained by a collection of Directory Service Agents (DSAs). The NSF is soliciting projects to deploy network information services, which will require significant efforts in organizing, populating and managing the Directory Information Base.

X.500 is not the only technology base that may be used in implementing network information services. Some techniques, such as Gopher[Alberti 91] and WAIS[Kahle 90] (Wide Area Information Service), will be used to retrieve information which is less directory oriented such as textual database retrieval and "how to?" type information.

3.4.3 Accounting

Public law 102-194 (formerly known as the S.272 Gore bill) mandates that the NREN shall have accounting mechanisms which allow users to be charged for usage of copyrighted materials available over the network and, where appropriate and technologically feasible, for usage of the network itself. An environment with well-established accounting mechanisms will be conducive to commercial entry into the market. The NREN will take advantage of accounting technologies which become available and will not attempt to develop NREN specific technologies. It is expected that related technologies will also directly impact traffic monitoring, analysis, modeling, and policy routing.

3.4.4 Network Tools for Hardening the Internet

The Internet today is a quilt of autonomous enterprises, each depending upon a slightly different set of technologies. Each network provider has different operational requirements and characteristics, such as 7 by 24 on call service and varying delivery of network uptime. One of the challenges to the NREN will be to provide higher reliability service, and quicker problem resolution. The current network users and operators have few and limited tools for monitoring the quality, behavior, and usage of the Internet and its component networks. This makes it difficult to take preemptive action in network maintenance and resource planning.

Currently, making the Internet "work", finding out what part is not working and subsequently fixing it, can be a tedious process involving numerous interactions between the users and the various providers of networking services. Successful operation of the NREN, and consequential user satisfaction, will depend on a mixture of services characterized automatic fault detection and isolation, and concomitant restoration of functionality even before a user sees a problem. This is the best of all possible worlds and could kick in at any level, from level 1 bitpipe to checkpointing routines in applications which can mask failures in network distributed applications.

The NSF will work with the NSFNET backbone providers, federal networks and operators of mid-level networks to jointly define operational criteria and aid in the development of tools to monitor conformance to these criteria.

3.5 Multiprotocol Support

As the Internet evolves from the confined IP environment to a large heterogeneous and global infrastructure it will be faced with the need to accommodate other packet level protocols, in particular protocols of the ISO. The OSI requirement is driven by the need to accommodate a large and international constituency, a U.S. Federal government requirement for future computer systems to be GOSIP compliant, as well as a potential way to address the shortcomings of the IP protocol suite, specifically the constraints of IP's 32 bit address space.

The Interagency Interim NREN will strategically adopt OSI CLNP to coexist with current IP services and is encouraging constituents to migrate from IP to CLNP over time. The time scale for IP and CLNP coexistence will be very long (greater than 5 years) and be driven by many factors, including technical requirements. It should be noted that this does not mandate a complete transition from the IP/TCP protocol stack to an OSI protocol stack. It is likely that TCP and UDP will run over CLNP, protecting the current investment in software which uses these protocols.

During the transition phase to CLNP, measures may have to be taken to circumvent IP address and routing limitations identified by the Internet community, as crystallized by the finding of the ROAD group.

4 Management Structure / Collaborative Support

The following groups currently form the committee structure tasked for planning and implementing the NREN within the context of the federal HPCC plan. This management structure is under periodic review and will be modified as required to adopt for changes in administrative requirements.

- The **Federal Network Council (FNC)** was formed in January 1990 as an interagency forum in support of the evolution of the Internet and other national research and education computer networks. This council is comprised of science research network project managers from all federal agencies involved in the HPCC. The **FNC** and its associated working groups provide for the management, coordination, and planning of the NREN program according to programmatic requirements.

The chairperson of the FNC and its Executive Committee act as liaisons to the relevant Executive Branch committees, including the the Office of Science and Technology Policy (OSTP), the Physical Mathematics and Engineering Sciences (PMES), the PMES High Performance Computing and Communications and Information Technology (HPCCIT) and its task forces, and the Office of Management and Budget (OMB). This relationship provides a mechanism to ensure that the NREN activities of FNC agencies are coordinated with the goals and requirements established by and for each of the participating agencies in conjunction with OSTP. NSF participates with the other NREN agencies on all of these committees.

- The NSF also participates fully in the FNC and its working groups, such as the **Engineering and Operations Working Group (EOWG)** , which is responsible for the design and implementation of the NREN, and other working groups related to policy, security and education.

- The **Federal Engineering and Planning Group (FEPG)** reports to the EOWG and is responsible for the technical analysis and operational coordination of the Federal agency NREN networks. Through participation with these various interagency committees and working groups, the NSF, in conjunction with the other agencies, strives to ensure a stable and vibrant networking infrastructure which satisfies its broad programmatic requirements in addition to the more focalized requirements of the HPCC NREN.

- In 1990, NSF chartered the **FNC Advisory Committee (FNCAC)** . Comprised of esteemed representatives from the telecommunications and computing industries, national laboratories and libraries, and academically-affiliated organizations, the FNCAC advises and informs the FNC on NREN issues.

5 Interagency Interim NREN Development Milestones

This section will identify NSF activities for FY92 which address the goals and requirements identified in this document. Many of these items involve engineering, implementation, and operation, and will therefore require coordination with appropriate Federal Network Council (FNC) working groups and affected NREN participants.

1. Gather, analyze, and identify possible solutions to the bandwidth, quality of service, connectivity, and support requirements of agencies, mid-levels, FARNET, libraries, and commercial service providers.

   NSF is currently working within the FNC committee structure to gather pertinent agency requirements. The NSF also sponsored the August 1991 FARNET meeting on Interregional Connectivity to gather requirements of the mid-level networks. NSF's National Science Board (NSB) presentation in November, 1991 identified the need for at least two service service providers and a routing arbiter to satisfy connectivity requirements identified by FARNET. FNC and IETF working groups are currently addressing other areas, such as the ROAD group efforts in the extension of Internet routing and addressing.

2. Identify and establish a mechanism whereby federal networks, commercial service providers, and mid-levels can effectively interconnect for the purpose of supporting the research and education community. The current FIX model will be generalized to a model based on Network Access Points (NAP). NAPs will be the focal point for routing coordination of federal, mid-level, international and other non-federal network service providers. The NAP model will accommodate IP and CLNP services and will interoperate with new level 2 services (e.g. DOE's SMDS trial).

   NSF will immediately assume responsibility for FIX East, currently operated by SURAnet at College Park, Maryland. The NSF engineering group (NEG) has a report in progress which details a FIX-East upgrade implementation plan to be completed by Summer 1992. NSF's active oversight for FIX-EAST should commence in June 1992 with the upgrade of FIX-East facilities and the initiation of a Network Access Point (NAP) beta test project at FIX-East.

3. Arrive at a formal set of criteria for mid-levels and other network service providers to directly connect to Network Access Points (NAPs).

   Requirements to be considered will include, but are not limited to, service quality, service offerings (IP, CLNP, BGP, IDRP, etc.), operations support, etc.

4. Develop a plan for the deployment of widespread network information services (NIS) serving the NSFNET and NREN community, evaluating X.500, WAIS, DNS, and other information search and retrieval capabilities to the NSFNET and NREN community.

   The utility of the current Internet is hampered by the lack of reasonable information about people, sites, and services within the Internet. The design and implementation of a distributed and delegated NIS system, including a NIS of first resort and a NIS of last resort, will ameliorate this situation. The NSF will issue a solicitation by April 1992 and make an award by the end of FY92, with deployment to commence in FY93.

5. Develop a plan accommodating the provisioning of interregional connectivity by more than one network service provider. The Internet is evolving from a central backbone model to a model consisting of a richly interconnected mesh infrastructure. The NAPs will be the cornerstone for the long term transition from a single generic transit backbone. A separate solicitation for both the new NSFNET "multiple service provider" backbone services and for the "routing arbiter/network access point manager" will be issued by June 1992, and awarded by April 1993 and become operational no later than April 94.

6. The NSF will formulate a transition plan with Merit, the backbone providers selected in the recompetition, federal agency networks and mid-level networks for the transition period from April 1993 to April 1994. The plan will focus on preserving a stable operating environment during the transition period.

7. Investigate Layer Two services through participation in DOE's Energy Sciences Network (ESnet) SMDS pilot project. The Switched Multi-Megabit Data Services (SMDS) testbed will aim for implementation by summer 1992, initially operating among five to seven ESnet sites at a minimum of 45 megabits per second. The project will also allow for the participation of other interested agencies, such as NASA and NSF. As a catalyst for increasing federal agency utilization of commercial network services, and the resulting incentive for greater deployment of these services by other carriers, this project is vital to the multi-agency NREN program and the HPCC initiative. In addition, this will allow the NSFNET to experiment with interconnecting dissimilar services such as T3 and SMDS technologies. NSF expects to subsidize a NSF site to participate in this project no later than the end of FY92.

8. Investigate and identify areas for enhancing the security of the network. Initially this will address technologies and techniques for for immediate deployment, at least in beta test mode. These include the deployment of one or more authentication schemes that can enhance the security of access to and remote use of supercomputer centers. This basis will then be used to advance the state and deployment of secure network management and remote network applications such as e-mail, file transfer, and remote login. The NSF will sponsor a security workshop no later than June 1992 in order to identify areas of NREN security to accelerate. The findings of this workshop will determine a specific schedule for further NREN and NSFNET security activities. NSF is currently working with the NIST and the FNC security working group to develop an FNC security plan, which

will later form part of the December 1992 OSTP report to Congress on the progress of the NREN.

9. In order to provide for a stable operational environment, encourage the development and dissemination of tools for network management and traffic monitoring, especially for higher speed networks. Also in pursuit of greater stability, develop a base set of variables to monitor on an operational basis, specific methods for inter NOC communications and problem solving, and an internetwork (among both backbones and mid-levels) trouble ticket coordination mechanism. In addition, identify a base service definition that interconnecting parties will subscribe to provide. NSF has initiated work in this arena in conjunction with relevant Federal agencies by sponsoring a FARNET "Tempering the Regionals" workshop in February 1992. The NSF will consult with the federal agency and mid-level networks, in concert with the findings from the FARNET meeting, to identify areas that can be addressed in NSF cooperative agreements.

10. Explore application layer peering, such as X.400 projects for interconnection of Commercial and Federal (Internet) OSI/GOSIP X.400 e-mail services, large scale global X.500 internetworking, as well as other applications such as "inter-realm authentication". The NSF will explore methods and encourage proposals for the purpose of advancing the state of a nationally integrated and interconnected directory service. In addition, NSF will also pursue means of enhancing the the state of service provider e-mail interconnectivity. Initial investigative actions in these areas will be concluded by June 1992 with the identification of any necessary near term actions or development.

6 Acknowledgements

Hans-Werner Braun and Peter Ford acknowledge support from the National Science Foundation for the study of engineering and architectural issues of the National Research and Education Network. Mr. Ford also acknowledges support from the Department of Energy Office of Energy Research (Contract No. KC0702) and Los Alamos National Laboratory.

References

[Alberti 91] Alberti et. al. "Notes on the Internet Gopher Protocol."" University of Minnesota, December 1991.

[Crowcroft 91] J. Crowcroft and I. Wakeman, *"Traffic Analysis of some UK-US Academic Network Data"*,
University College London, Department of Computer Science, March, 1991.

[Danzig 92] Danzig, Peter B., Obraczka, Katia, and Kumar, Anant. *An Analysis of Wide-Area Name Server Traffic, submitted to Sigcomm 92*

[FEPG 91] Federal Networking Council (FNC) Engineering and Planning Group, *"Architectural Requirements for the National Research and Education Network, Version 1.0"*,
January 22, 1991.

[FRICC] Federal Research Internet Coordinating Committee, *"Program Plan for the National Research and Education Network"*,
May 23, 1989.

[HWB 92] Hans-Werner Braun for the NEG, *"Redesign and Extension of Multi-Agency Interconnection Points* ,
March 12, 1992.

[ISO 9594-1] Information Processing Systems – Open Systems Interconnection – The Directory – Overview of Concepts, Models, and Service. International Electrotechnical Committee, December, 1988. International Standard 9594-1.

[Kahle 90] Kahle, Brewster. "WAIS Interface Prototype Functional Specification." Thinking Machines Corporation, April 1990.

[Lottor 92] Lottor, M. "Internet Growth (1981-1991)", RFC 1296, January 1992.

[Lougheed 91] Lougheed, K., Rekhter, Y., "A Border Gateway Protocol", RFC1167, October 1991.

[PMES 91] Federal Coordinating Council for Science, Engineering, and Technology, *"Grand Challenges: High Performance Computing and Communications"*,
Supplement to the President's Fiscal Year 1992 Budget.

[PMES 92] Committee on Physical, Mathematical, and Engineering Sciences; Federal Coordinating Council for Science, Engineering, and Technology, and; Office of Science and Technology Policy, *"Grand Challenges: High*

Performance Computing and Communications - The FY 1992 U. S. Research and Development Program",

[Rekhter 89] Rekhter, J., "EGP and Policy Based Routing in the New NSFNET Backbone", RFC1092, February 1989.

A Acceptable Use

Some network service providers are bound to rules governing network usage, sometimes even requiring official Acceptable Use Policies. Other providers, especially on the commercial level, transmit any traffic that does not break applicable laws. Legal standards of behavior for transit providers have not yet been clearly delineated, creating problems in determining the extent of liability of a network provide transporting traffic of questionable legal or moral nature. While the regulated phone companies are indemnified from liability for the use of their network, the FCC currently has not issued decisions regarding data networks. Please see 6 for a presentation of NSF's current acceptable use policy.

Complicated transit policies in federal, corporate, and international realms will all require the accommodation of specific high performance, high bandwidth needs, as well as the aggregation of traffic from many users. Although it may expand the work scope of network management, emerging network technology must allow for the implementation of transit policies of service providers, including those of the federal government.

A.1 Current NSF Acceptable Use Policy

The NSFNET Backbone Services Acceptable Use Policy is derived from the previous "interim" policy under the advice of the Networking Division's Advisory Committee and the NSF General Council. In the NSF General Council's opinion, this represents the most liberal possible policy consistent with the NSF enabling legislation, as amended. Compared with the Interim policy, the most obvious liberalizations are items 8 and 11.

General Principle

1. NSFNET Backbone services are provided to support open research plus research arms of for-profit firms when engaged in open scholarly communication and research. Use for other purposes is not acceptable.

Specifically Acceptable Uses

2. Communication with foreign researchers and educators in connection with research or instruction, as long as any network that the foreign user employs for such communication provides reciprocal access to US researchers and educators.

3. Communication and exchange for professional development, to maintain currency, or to debate issues in a field or subfield of knowledge.

4. Use for disciplinary-society, university-association, government-advisory, or standards activities related to the user's research and instructional activities.

5. Use in applying for or administering grants or contracts for research or instruction, but not for other fundraising or public relations activities.
6. Any other administrative communications or activities in direct support of research and instruction.
7. Announcements of new products or services for use in research or instruction, but not advertising of any kind.
8. Any traffic originating from a network of another member agency of the Federal Networking Council if the traffic meets the acceptable use policy of that agency.
9. Communication incidental to otherwise acceptable use, except for illegal or specifically unacceptable use.

Unacceptable Uses

10. Use for for-profit activities (consulting for pay, sales or administration of campus stores, sale of tickets to sports events, and so on), or use by for-profit institutions unless covered by the General Principle or as a specifically acceptable use.
11. Extensive use for private or personal business.

This statement applies to use of the the NSFNET Backbone only. NSF expects that connecting networks will formulate their own use policies. The NSF Division of Networking and Communications Research and Infrastructure will resolve any questions about this Policy or its interpretation.