Redesign and Extension of Multi-Agency Interconnection Points

Hans-Werner Braun for the NEG

D R A F T

June 1, 1992

This document is a result of several discussions regarding the current Federal IntereXchange (FIX) environment. Particular refinements have resulted from conversations with Peter Ford (LANL/NEG), Bob Aiken (NSF), Milo Medin (NASA), Tony Hain (LLNL/DOE) and Claudio Topolcic (CNRI).

Abstract

In its role to support the NREN effort, the NSF has asked the NEG Engineering Group (NEG) to draft a document describing the potential redesign and architectural extension of one of the interconnection points for national US agency networking activities. This document addresses in particular FIX-East, located near the University of Maryland, but we hope the issues and suggestions are applied to the larger NREN community, in particular in the context of generic access instrumentation for the Internet.

Contents

1 Introduction

Over the last few years, the US federal agencies which are involved in Internetworking have recognized the need to directly interconnect their backbones via a well defined architecture. The mechanisms developed have been largely ad hoc, with no guiding architectural principles in design. This approach is surprising given the fairly rigid set of constraints which such an interconnection architecture need satisfy. The NSF has asked the NREN Engineering Group (NEG) to investigate a potential redesign and architectural extension of one of the interconnection points for federal US agency networking activities. This document describes the national R&E requirements of this architecture, and in the process offers a set of definitions which we hope will clarify possible ambiguities of the current Internet environment.

This multiple entity interconnection architecture (FIX) needs to support a largely production oriented environment, maintaining reliable and stable networking access for scientists, researchers, and educators. Federal agencies have traditionally maintained their own independent networks, which is not surprising given their diverse, albeit sometimes complementary, goals and implementation technologies. Indeed, a single shared infrastructure is neither a necessary nor sufficient condition for harmonizing multiple agency networks, and the NREN architecture certainly does not by definition require it. A disjoint conglomeration of multiple independent providers can perform the same functionality, but such an infrastructure further amplifies the need for a robust interconnection design.

The continuing and rather evolving importance of interconnection architectures, both in terms of services and technologies, requires a clear definition of goals, requirements, and expectations for the community, both from the user and network perspective. Integrity of the infrastructure will become even more critical in the near future, as the Internet must accommodate an even greater number and variety of service providers. Thus, we also hope to take this opportunity to extend and continuously evolve the fundamental notion of a network interconnection point. We offer criteria for such locations, encompassing the gatewaying of services at multiple levels, including at the application level, and allowing for a graceful integration of resources such as statistics gathering, traffic monitoring, and servers for time and root domain name services. We hope to address an obvious need in the community today: a clear outline of terms and conduct for interconnection activities.

2 Current Interconnection Strategy

The United States component of the Internet is based on a three-level hierarchy of national scope: federal agency R&E backbones, attached mid-level networks, and connected local sites (Figure 2). We briefly examine the topology as a logical tree rooted at the two existing Federal Interexchange points (FIXes), with branches representing the federal agency networks. This abstraction is not wholly accurate, as it ignores commercial IP providers, international links, and bypass interconnections. Nonetheless, the structural description provides an architectural understanding for how networks are interconnected and local responsibility compartmentalized.

U.S. federal agency backbone interconnection currently occurs through two specific points, where routers are co-located on a local area network with robust operational support. These two IntereXchange points are: FIX-West (Figure ??), at NASA-Ames at Moffet Field, CA, and FIX-East, at a SURANET location outside of the University of Maryland at College Park. The Federal Engineering Planning Group principally designed these backbone interconnection points, and initially implemented FIX-West. The second interexchange point, FIX-East, has been redesigned and partly moved as a collaborative multiagency effort since its establishment, with much of the changes having been driven by Milnet requirements.

3 Current FIX-West topology

Networks attached to these FIXes form autonomously managed Administrative Domains (AD), which exchange routing information by means of interdomain protocols, typically EGP. These protocols allow networks to selectively maintain peer relationships by controlling the flow of routing information between the agency networks. Route filtering, among other mechanisms in standard Internet protocols, allows the definition of a logical topology of interconnected federal agency backbones and mid-level networks. Using such

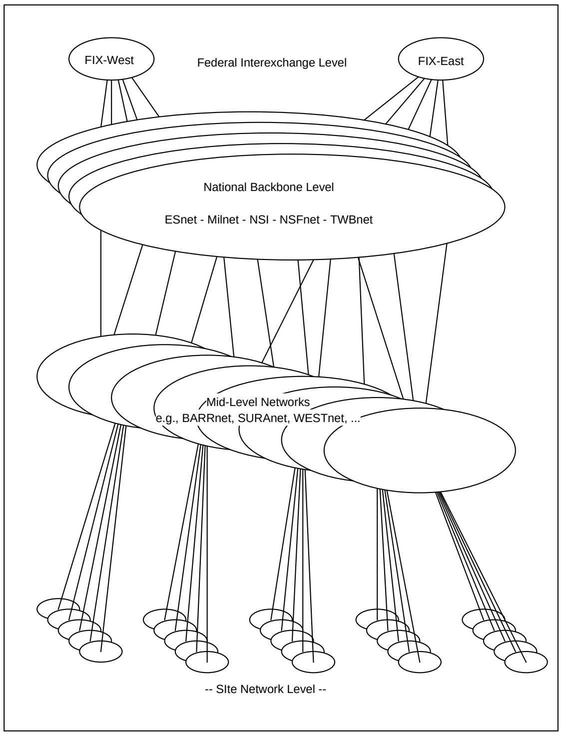

Figure 1: Hierarchical Network Structure

selective propagation schemes and preferred routes at the FIXes, with their strategic location and cross-section of connecitivity, facilitates the robust implementation of interagency transit policies.

Unfortunately, selective manual isolation of routing information inhibits the sharing of available bandwidth and eliminates the ability to use redundant connectivity to repair network partitions which may occur due to backbone network failures. An evolution in the concept of interconnection points is clearly essential to robust connectivity and well-defined and implemented routing policies in today's ever-changing environment.

NSF has recently issued a contract for the management of the FIX-East network interconnection point. The NEG would like to take this opportunity to use the FIX-East design to represent a prototype of a more general entity: a Network Access Point (NAP). The two major criteria for redesign of FIX-East are: one, to more closely parallel the FIX-West interconnection architecture and functionality; and two, to extend FIX-East to allow for coordinated and controlled connection to non-federal service providers. At the same time, it is important that the design allow for the preservation of an agency viewpoint, and satisfy agency objectives at those sites which house and support an interconnection point. The fact that the performance of the interconnection points often involves issues of national security further magnifies this requirement.

4 Interconnection Paradigms

The current and somewhat restrictive definition of a FIX, to interconnect only federal agency R&E networks, attempts to provide a manageable, dependable and predictable environment. Such a stable atmosphere is critical to all FIX clients; thus an evolution from the current scheme must not result in any reduction of performance or services. However, the dramatic expansion in scope of the infrastructure, makes some evolution of the current top level federal US portion of the Internet absolutely essential. New service providers, who often offer national and comprehensive services, consider themselves entitled to equal rights

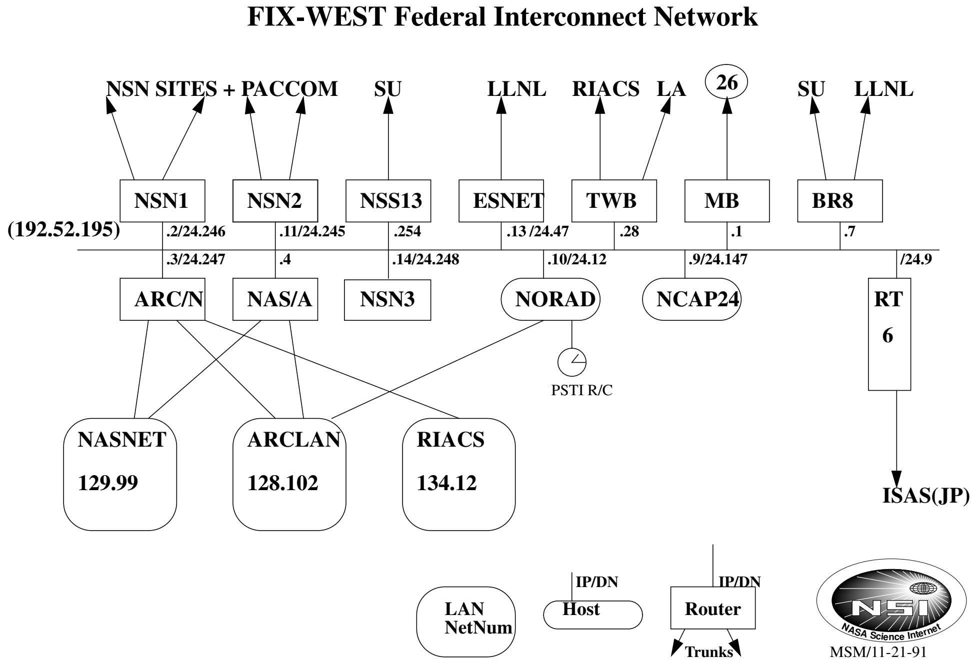

Figure 2: FIX-West Topology

to “fair access” to connectivity into the high-level network hierarchy.

Agency networks (e.g., NSI, TWBnet, ESnet) define agency-related traffic as their objective In contrast, the NSFNET backbone, and its associated mid-level networks, assume responsibility for transmission of any R&E traffic meeting the NSFNET Appropriate Use Policy (see Appendix A). 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, in its role as provider of the current NSFNET backbone, acts as “routing arbiter”, historically providing all networks with a picture of Internet routing topology, and enforcing the NSF AUP. This routing arbiter function has gradually increased in complexity 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.

In fact, the drive for commercial network interconnectivity has led to the formation of commercial interexchange points (CIXes), analogous to the FIX concept for commercial backbones and mid-levels supporting non-AUP compliant traffic. Unfortunately, the increasingly convoluted topology has exposed several weaknesses in the Internet routing technology base, which was not designed with the NSF AUP in mind. Private vendors will gradually assume even more of the now federally-subsidized market for wide area internetworking services and infrastructure. Mid-level, federal agency, and corporate customers will in turn increase their contractual relationships with such commerical providers of both inter-regional connectivity, and database and information services.

Greater competition in the "Internet industry" will lead to a broader set of viable alternatives for network clients, 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.

NSFNET itself, in its role as the general purpose transit network for generic R&E traffic will also sollaborate with commercial service providers, creating a complex web of interrelation between Internet component networks. Distributing the transit function across multiple networks controlled by different administrations poses severe implementation and management difficulties.

The current FIX interconnection points illustrate this point: in reality, they have already developed into richly connected environments supporting agency, regional, and commercial connectivity. They have suffered, however, from the lack of a concerted engineering effort toward a rational implementation of this connectivity. As the Internet community grows in size and variety, including non-R&D participants, it becomes more difficult, and critical, to provide well-designed and open access for new providers, while realistically addressing the goals and requirements of all participants. We contend that a carefully designed evolution of the current FIX architecture can minimize difficulties in an otherwise problematic environment.

5 Network Access Points

5.1 Order Into Chaos

Current routing technologies work best in a pure tree hierarchy, which is not an appropriate topological model given the scope, scale and interconnectivity of current national network providers.

The developing Border Gateway Protocol, arose from the perceived need to evolve away from this strict model toward a more dynamic topology. For overall routing manageability, BGP may prove superior to mechanisms such as strict route filtering based on Autonomous System identifiers for a whole AS path. Wide scale deployment of BGP will facilitate migration of the Internet routing topology over the next few years. Evolution of the current logical interconnection architecture from an abstract tree toward a mesh of multiple connections at disparate service qualities may lead to asymmetric routing and transit paths. Furthermore, the transition will not occur overnight; the NEG expects a two year lag at least. Thus an interconnection scheme must temporarily support both architectures and even translate between them. NSF will fund and manage such a scheme, with an aim toward a future scenario where the participating networks fund and manage the NAPs.

5.2 Requirements of the Job

Any viable interconnection strategy must allow Internet network providers, including non-federal networks, reasonable access to the Interagency Interim NREN. Maintaining a coherent picture of routing within the Internet, currently the responsibility the closed federal R&E community at the FIXes, should diversify toward a system where all network providers participate in shaping the routing fabric at **Network Access Points (NAPs)**. NAPs must be robust in their accommodation of the following diverse requirements:

1. Agencies require disambiguation from non-federal networks. This may include physical or logical separation from non-federal networks. A router interfacing the non-federal service providers may effect a physical separation. Route servers which can collapse routing information of multiple service providers into a single and coherent view may function as logical separators. A NAP route server may even function as a dedicated administrative domain, using a single Autonomous System number for their routing advertisement.

2. Agencies may have a requirement to attech to non-federal networks, in addition to needing direct access to a service provider, such as a mid-level network or some commercial enterprise. Specific service provisioning of an agency network could mandate such direct connection.

3. Route servers must provide client networks with a consistent set of information, if necessary by sanitizing routing information from multiple service providers. The relation between the administrative domains must be codified, at least in terms of firewalling, routing configurations, acceptable topologies and usage policies.
4. NAP attachment points should have a minimal set of well-defined and published standard requirements and procedures for the accomodation of new clients. These procedures should include a definition of appropriate routing protocols and a standardized methodology for acceptable routing and transit of packets among client networks. Selected protocols should simplify peering relationships between client networks, reducing the need for a full matrix of inter-domain routing relationships. All networks connected to NAP peering points would be expected to subscribe to this base set of operational requirements.
5. NAPs must allow for an integration of multiple usage policies, including those of the federal and non-federal components. The NAP operations contractor, as well as a NAP oversight group that NSF will charter, must have all usage policy documents available.
6. Instrumentation for statistics gathering and inter-provider accounting.
7. Other servers, including those for application level gateways. The extended interconnection environment must accomodate a variety of other services, such as root domain name servers, network time, monitoring and statistics, and perhaps top level future X.400 and X.500 gateway services. Dedicated instrumentation connected to either a specific or a jointly available LAN, could provide such services. Operationally feasible multilayer interconnections will require careful design and a well-understood architecture.

The management and operation of an intermediate router or route server must be under the control of the participating agencies. The National Science Foundation, in its infrastructure coordination role, should take the lead in ensuring the interconnection of multiple service providers. NSF should establish management structures such that the community can clearly understand and effectively interact with the day-to-day NAP operators, through well-defined communication and management processes for addressing problems. NSF should immediately establish a NAP oversight group that provides NAP operators with guidance on NAP design and management. A local service provider should understand that their rights to the NAP, particularly the federal FIX portion, are limited, and that they assume several responsibilities in return for the operations contract:

- The NAP services provider is not in a position to modify the FIX component of the NAP in any way without explicit permission of the federal agencies via the NAP oversight group.
- The provider is responsible for local security and must provide UPS power and local operations support on a 24/7 basis.
- The provider must restrict facility access to authorized or escorted people.
- The provider is advised to participate in discussions with participants of the FIX and NAP connected networks to ensure robust services.

The operations contractor must provide a service guarantee with respect to management, operation and administration. In order to avoid excessive failure points, the LAN connections must serve as concentrators for the FDDI and Ethernet connections. These LANs can then accomodate routers of the respective peer groups. The performance of the LANs should be consistently monitored, and NAP oversight groups should receive relevant capacity-planning information in a timely manner.

In order to preserve the current FIX architecture, only bona fide agency network participants would be allowed to connect to the FIX component of the NAP. [1] Another set of allocated interfaces would provide more general access, constituting a General InterExchange (GIX). The GIX and the FIX together consitute a Network Access Point (NAP), allowing federal and non-federal networks to gracefully coexist. 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.

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

6 Redesign and Extension

We now address the specific case of FIX-East and its potential as a NAP prototype.

As described above, a first major step in the evolution of FIX-East is to gracefully accommodate non-federal service providers. In pursuit of this objective, the NSF should initially implement an east-coast NAP with multiple external interfaces, including both Ethernet and FDDI. The equipment must be fast enough to accomodate all the performance requirements for NAP interconnectivity. These requirements depend on the number and performance of the client networks, and must be periodically reevaluated by the NAP oversight group in conjunction with the NAP operations contractor.

The NEG offered one suggested initial NAP-East design, illustrated in figure 7.

7 Proposed NAP-East topology

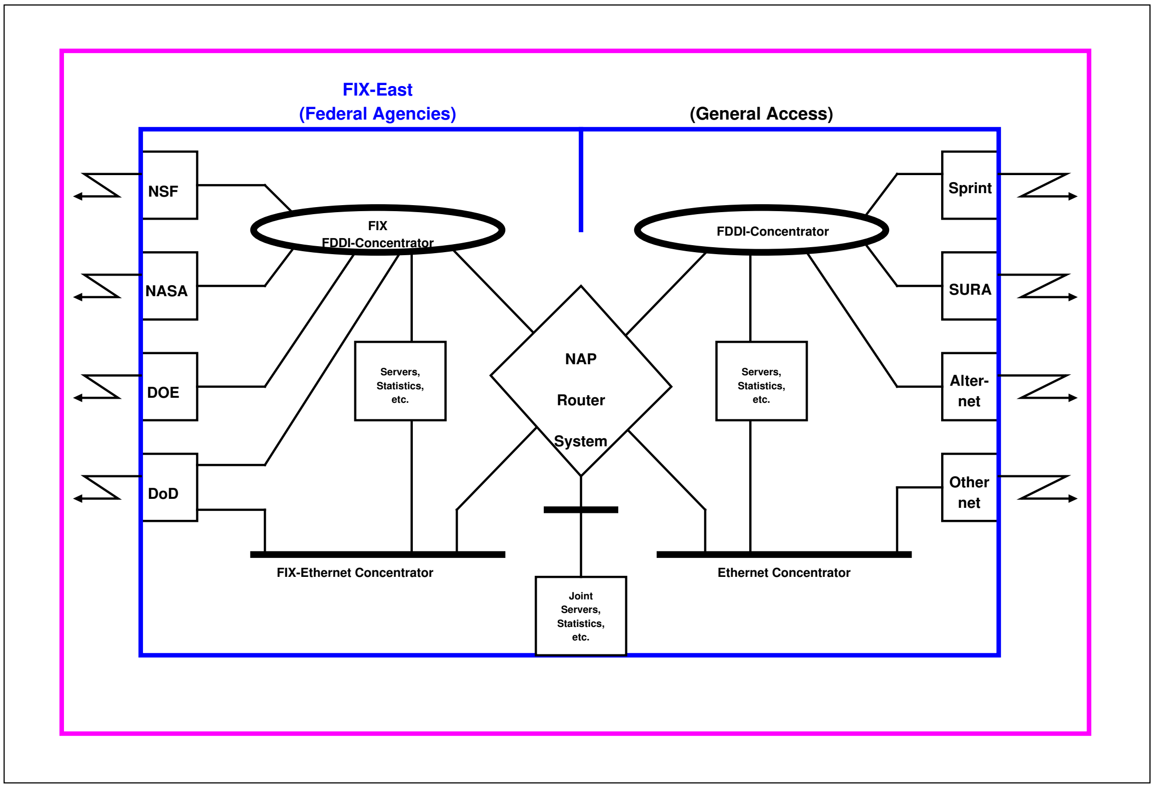

Figure 3: NAP Design, Version 1

As a result of DOE considerations for the ability to interconnect in a Level 2 environment, Tony Hain has suggested an alternative NAP-East implementation, shown in figure 7. This latter configuration interfaces all the routers directly via a Level 2 connection, rather than using a Level 3 router. Hain's approach still requires a dedicated NAP routing server to accommodate agencies who want to more strongly decouple themselves from non-federal routers.

From FEPG discussions, NASA (represented in this case by Milo Medin) disagrees with this approach. Medin claims that with no intermediating router, this configuration does not allow for an easy disambiguation of a federal and a non-federal portion of client networks. NASA would like to see instead a single interface with consistent routing information offered to non-federal service providers. An intermediate router, capable of aggregating routes from multiple clients, filtering the information, and subsequently redistributing it to routers, could accomplish this task.

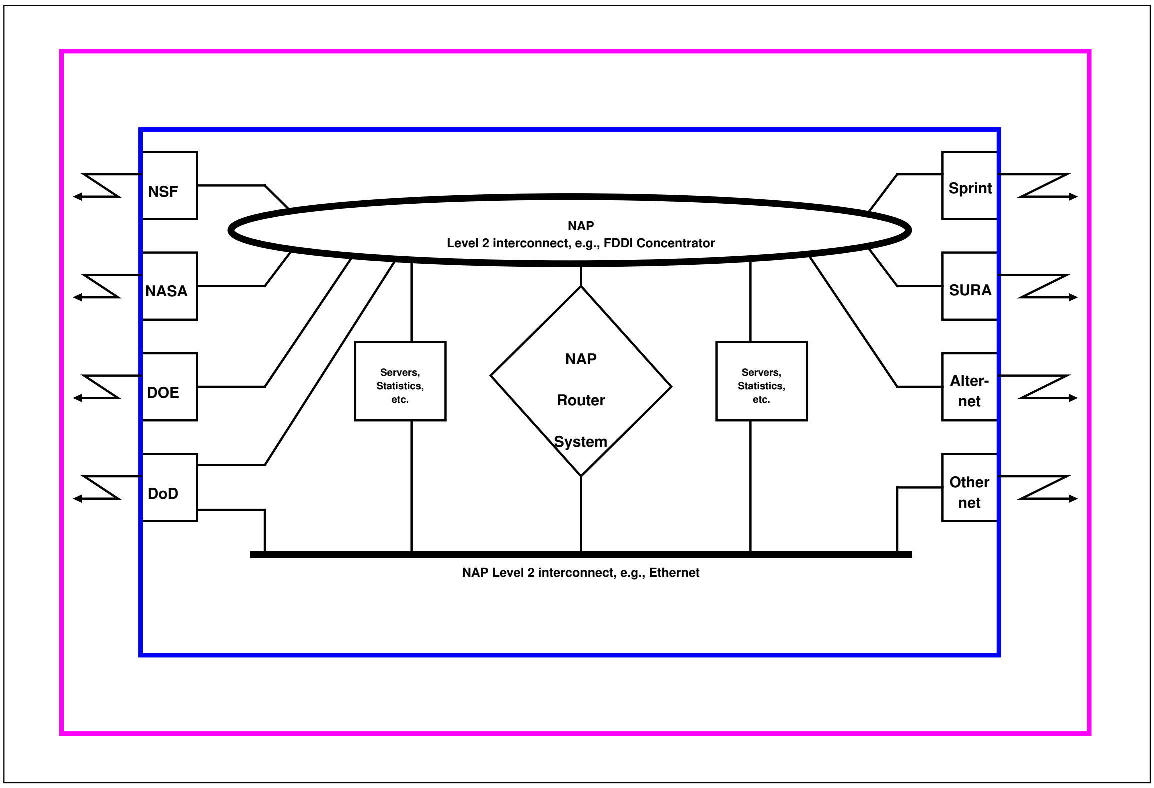

Figure 4: NAP Design, Version 2

The functionality of both scenarios can be achieved with a router to specifically interface between the federal and non-federal LANs, while allowing agency service providers to connect to any LANs they choose. Connecting to multiple LANs would, in terms of hardware, require only an additional interface card for the already existing federal router, after the necessary advance routing coordination.

8 Conclusion: The Future of NAPs

NAPs will require robust implementations, as the interconnected networking environment will heavily depend on their architecture. Important components of robustness include multiple, geographically dispersed NAP points, operational support and stability of the physical environment, and a stable architecture for routing between the various administrative domains.

As the Internet environment extends to a greater number and variety of service providers, and a single shared infrastructure becomes a less realistic scenario, interconnection becomes a contentious issue. An architecture for the graceful accommodation of new providers is essential to dispel the polarization between the commercial and federal communities which exists today.

In pursuit of this goal, we offer this document as more than a discussion of the current situation at FIX-East. We hope we have also provided a fundamental set of acceptable definitions which will allow intelligent discussion in the community concerning the appropriate evolutionary path away from the current notions of FIXes and CIXes to a more comprehensive architecture. Increasing the efficacy of critical interconnection activities will enhance the productivity and usefulness of the Internet as a whole.

8.1 Timelines

- Receive funding proposal from SURAnet

- Proposal review finished
- Funding award
- Implement NAP-East router
  1. Two FDDI interfaces
  2. Two Ethernet interfaces
- Implement FDDI and Ethernet concentrators
- Management control of NAP-East router to NSF
- Implement routing control in the NAP-East router
- Move agency routers to the new FIX-East
- Integrate GIX services
- Implement Servers for Statistics and Time
- Implement application level gateways (X.400, X.500, DNS, etc.)

8.2 Cost Factors

- NAP-East router, including 2 * FDDI and 2 * Ethernet
- Concentrators for two FDDI rings and two Ethernets
- FDDI and Ethernet cables
- Workstations, one each for FIX and GIX management/statistics/etc.
- SURAnet operations support
- NSFNET management support