

Implementation Plan for Interagency Research Internet

STATUS OF THIS MEMO

The RFC proposes an Interagency Research Internet as the natural outgrowth of the current Internet. This is an "idea paper" and discussion is strongly encouraged. Distribution of this memo is unlimited.

OVERVIEW

Networking has become widespread in the scientific community, and even more so in the computer science community. There are networks being supported by a number of the Federal agencies interested in scientific research, and many scientists throughout the country have access to one or more of these networks. Furthermore, there are many resources (such as supercomputers) that are accessible via these networks.

While many of these networks are interconnected on an informal basis, there is currently no consistent mechanism to allow sharing of the networking resources. Recognizing this problem, the FCCSET Committee on Very High Performance Computing formed a Network Working Group. This group has recommended an administrative and management structure for interconnecting the current and planned agency networks supporting research. The structure is based on the concept of a network of networks using standard networking protocols.

This report elaborates on the earlier recommendation and provides an implementation plan. It addresses three major areas; communications infrastructure, user support, and ongoing research. A management and administrative structure is recommended for each area, and a budgetary estimate provided. A phased approach for implementation is suggested that will quickly provide interconnection and lead to the full performance and functionality as the required technologies are developed and installed. While this report addresses the interconnection of agency networks, and cooperation by certain federal agencies, some discussion is presented of the possible role that industry can play in support and use of such a network.

Work reported herein was supported by Cooperative Agreement NCC 2-387 from the National Aeronautics and Space Administration (NASA) to

the Universities Space Research Association (USRA). This report was prepared in response to a request from John Cavallini, Chairman of the Networking Working Group of the FCCSET Committee on Very High Performance Computing.

INTRODUCTION

Computer networks are critical in providing scientists access to computing resources (such as supercomputers) and permitting computer supported interaction between researchers. Several agencies, recognizing this need, have established networks to provide the needed communications infrastructure. The need for this infrastructure, though, cuts across the various agencies. To that end, the FCCSET Committee on Very High Performance Computing Network Working Group has recommended the formation of an Interagency Research Internet (IRI) [1].

The purpose of this report is to suggest an implementation plan for such an IRI. It addresses three major areas; communications infrastructure, user support, and ongoing research. A management and administrative structure is recommended for each area, and a budgetary estimate provided. A phased approach for implementation is suggested that will quickly provide interconnection and lead to the full performance and functionality as the required technologies are developed and installed. Finally, some discussion is presented on a possible role for industry in supporting and using such a network.

Motivation

The prime responsibility for providing the required infrastructure for successful research lies with the researcher, his/her institution, and the agency supporting that research. Thus, the individual agencies have installed and are continuing to enhance computer networks to allow their researchers to access advanced computing resources such as supercomputers as well as being able to communicate with each other via such facilities as electronic mail.

However, there are a number of reasons why it is advantageous to interconnect the various agency networks in a coherent manner so as to provide a common "virtual" network supporting research.

The need to make effective use of available networks without unnecessary duplication. The agencies each support researchers in many parts of the country, and have installed equally widespread resources. Often, it is more effective for a scientist to be provided networking service through a different agency network than the one funding his research. For example, suppose several scientists at an institution are already being funded by NASA and

are connected to a NASA supported network. Now a scientist at the same institution but supported by NSF needs access to an NSF supercomputer. It is much more effective to provide that connectivity through an interconnection of NASA and NSF networks than to establish another connection (to NSFnet) to the same university.

The need to establish communication infrastructure to permit scientists to access resources without regard to which network they are connected but without violating access controls on either the networks or the resources. A scientist may be supported by multiple agencies, and therefore have access to resources provided by several agencies. It is not cost-effective to have to provide a separate network connection to the scientist for each of those agency resources.

The need for a communications infrastructure to encourage collaborative scientific research. One of the primary functions of a computer network supporting science is the encouraging of collaboration between researchers. Scientific disciplines typically cut across many different agencies. Thus, support of this collaboration should be without regard to agency affiliation or support of the scientists involved.

The need for a cooperative research and development program to evolve and enhance the IRI and its components where appropriate. Scientific research is highly demanding of both the computing and networking environment. To assure that these needs continue to be met, it is necessary to continually advance the state of the art in networking, and apply the results to the research networks. No individual agency can afford to support the required research alone, nor is it desirable to have inordinate duplication of research.

Summary of previous report

These reasons led to the formation of the FCCSET Committee on Very High Performance Computing and its Network Working Group. This group began in early 1985 to discuss the possibility of interconnecting into a common networking facility the various agency networks supporting scientific research. These discussions led to the report issued earlier this year [1] recommending such an approach.

The report used the "Network of Networks" or Internet model of interconnection. Using a standard set of protocols, the various networks can be connected to provide a common set of user services across heterogenous networks and heterogenous host computers [2, 3,4]. This approach is discussed further in the Background section

below.

The report goes on to recommend an administrative and management structure that matches the technical approach. Each agency would continue to manage and administer its individual networks. An interagency body would provide direction to a selected organization who would provide the management and operation of the interconnections of the networks and the common user services provided over the network. This selected organization would also provide for coordination of research activities, needed developments, and reflecting research community requirements into the national and international standards activities.

Overview of Implementation Plan

The general structure of the proposed IRI is analogous to a federal approach. Each of the agencies is responsible for operating its own networks and satisfying its users' requirements. The IRI provides the interconnecting infrastructure to permit the users on one network to access resources or users on other networks. The IRI also provides a set of standards and services which the individual agencies, networks, and user communities can exploit in providing capabilities to their individual users. The management structure, likewise, provides a mechanism by which the individual agencies can cooperate without interfering with the agencies' individual authorities or responsibilities.

In this report, an implementation plan for the IRI is proposed. First, some background is given of the previous efforts to provide networks in support of research, and the genesis of those networks. A description of the suggested approach to attaining an IRI is then given. This description is divided into two sections; technical and management. The technical approach consists of two components. First is the provision of an underlying communications infrastructure; i.e. a means for providing connectivity between the various computers and workstations. Second is provision of the means for users to make effective use of that infrastructure in support of their research.

The management section elaborates on the suggestions made in the FCCSET committee report. A structure is suggested that allows the various agencies to cooperate in the operations, maintenance, engineering, and research activities required for the IRI. This structure also provides the necessary mechanisms for the scientific research community to provide input with respect to requirements and approaches.

Finally, a phased implementation plan is presented which would allow

the IRI to be put in place rapidly with modest funding. A budgetary estimate is also provided.

BACKGROUND

The combination of packet switched computer networks, internetworking to allow heterogeneous computers to communicate over heterogeneous networks, the widespread use of local area networks, and the availability of workstations and supercomputers has given rise to the opportunity to provide greatly improved computing capabilities to science and engineering. This is the major motivation behind the IRI.

History of Research Network

The Defense Advanced Research Projects Agency (DARPA) developed the concept of packet switching beginning in the mid 1960's. Beginning with the Arpanet (the world's first packet switched network) [5], a number of networks have been developed. These have included packet satellite networks [6,7], packet radio networks [8,7], and local area networks [9].

Although the original motivation for the Arpanet development was computer resource sharing, it was apparent early on that a major use of such networks would be for access to computer resources and interaction between users [10]. Following the Arpanet development, a number of other networks have been developed and used to provide both of these functions [11]. CSNET was initiated to provide communications between computer science researchers [12,13]. CSNET was initiated by the NSF in cooperation with a number of universities, but is now self-sufficient. Its subscribers include universities throughout the world as well as industrial members interested in interacting with computer scientists.

CSNET makes use of a number of networking technologies including the Arpanet, public X.25 networks, and dial-up connections over phone lines, to support electronic mail and other networking functions. In addition to the basic data transport service, CSNET and Arpanet operate network information centers which provide help to users of the network as well as a number of services including a listing of users with their mail addresses (white pages) and a repository where relevant documents are stored and can be retrieved.

With the installation of supercomputers came the desire to provide network access for researchers. One of the early networks to provide this capability was MFEnet [11]. It was established in the early 1970's to provide DOE-supported users access to supercomputers, particularly a Cray 1 at Lawrence Livermore National

Labs. Because MFEnet was established prior to widespread adoption of the TCP/IP protocol suite (to be discussed below), the MFEnet uses a different set of protocols. However, interfaces have been developed between the MFEnet and other networks, and a migration plan is currently under development.

NASA Ames Research Center has long been in the forefront of using advanced computers to support scientific research. The latest computing facility, the Numerical Aerodynamic Simulator, uses a Cray 2 and other machines along with a number of networking technologies to provide support to computational fluid dynamics researchers [14]. This system uses the TCP/IP protocol suite both locally and remotely and provides easy access through advanced workstations.

Recognizing the importance of advanced computers in carrying out scientific research, NSF in 1984 embarked on an ambitious program to provide supercomputer access to researchers. This program involved both the provision of supercomputers themselves (through purchase of computer time initially, and establishment of supercomputer centers) and provision of access to those supercomputers through an extensive networking program, NSFnet [15]. The NSFnet uses a number of existing networks (e.g. Arpanet, BITNET, MFEnet) and exploratory networks interconnected using the TCP/IP protocol suite (discussed below) to permit scientists widespread access to the supercomputer centers and each other. The NSFnet is also taking advantage of the widespread installation of campus and regional networks to achieve this connectivity in a cost effective manner.

The above are only a small number of the current and existing networks being used to support research. Quarterman [11] provides a good synopsis of the networks currently in operation. It is obvious from this that effective interconnection of the networks can provide cost-efficient and reliable services.

Starting in the early 1970's, recognizing that the military had a need to interconnect various networks (such as packet radio for mobile operation with long-line networks like the Arpanet), DARPA initiated the development of the internet technologies [16]. Beginning with the development of the protocols for interconnection and reliable transport (TCP/IP), the program has developed methods for providing electronic mail, remote login, file transfer and similar functions between differing computers over dissimilar networks [4,3]. Today, using that technology, thousands of computers are able to communicate with each other over a "virtual network" of approximately 200 networks using a common set of protocols. The concepts developed are being used in the reference model and protocols of the Open Systems Interconnection model being developed by the International Standards Organization (ISO) [17].

This is becoming even more important with the widespread use of local area networks. As institutions install their own networks, and need to establish communications with computers at other sites, it is important to have a common set of protocols and a means for interconnecting the local networks to wide area networks.

Internet Model

The DARPA Internet system uses a naming and addressing protocol, called the Internet Protocol (IP), to interconnect networks into a single virtual network. Figure 1 shows the interconnection of a variety of networks into the Internet system. The naming and addressing structure allows any computer on any network to address in a uniform manner any computer on any other network. Special processors, called Gateways, are installed at the interfaces between two or more networks and provide both routing amongst the various networks as well as the appropriate translation from internet addresses to the address required for the attached networks. Thus, packets of data can flow between computers on the internet.

Because of the possibility of packet loss or errors, the Transmission Control Protocol (TCP) is used above the IP to provide for reliability and sequencing. TCP together with IP and the various networks and gateways then provides for reliable and ordered delivery of data between computers. A variety of functions can use this connection to provide service to the users. A summary of the functions provided by the current internet system is given in [4].

To assure interoperability between military users of the system, the Office of the Secretary of Defense mandated the use of the TCP/IP protocol suite wherever there is a need for interoperable packet switched communications. This led to the standardization of the protocols [18, 19, 20, 21, 22].

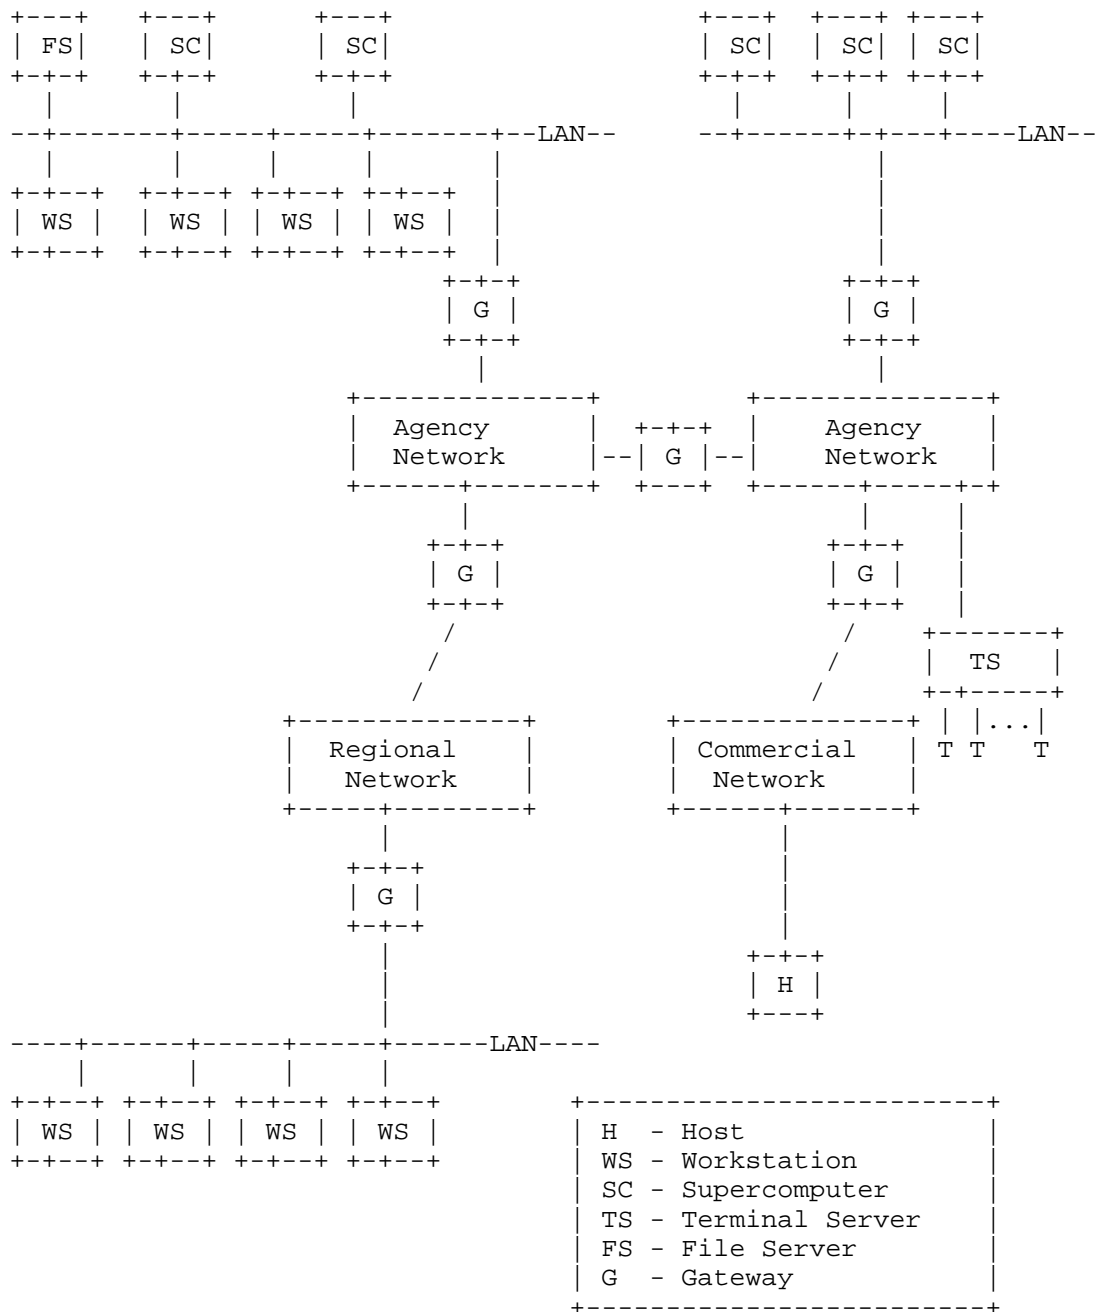


Figure 1: Internet System

Thus, the TCP/IP protocol suite and associated mechanisms (e.g. gateways) provides a way to interconnect heterogeneous computers on heterogeneous networks. Routing and addressing functions are taken care of automatically and transparently to the users. The ISO is currently developing a set of standards for interconnection which are very similar in function to the DARPA developed technologies. Although ISO is making great strides, and the National Bureau of Standards is working with a set of manufacturers to develop and demonstrate these standards, the TCP/IP protocol suite still represents the most available and tested technology for interconnection of computers and networks. It is for that reason that several agencies/programs, including the Department of Defense, NSF and NASA/NAS, have all adopted the TCP/IP suite as the most viable set of standards currently. As the international standards mature, and products supporting them appear, it can be expected that the various networks will switch to using those standards.

TECHNICAL APPROACH

The Internet technology described above provides the basis for interconnection of the various agency networks. The means to interconnect must satisfy a number of constraints if it is to be viable in a multi-agency environment.

Each agency must retain control of its own networks. Networks have been established to support agency-specific missions as well as general computer communications within the agency and its contractors. To assure that these missions continue to be supported appropriately, as well as assure appropriate accountability for the network operation, the mechanism for interconnection must not prevent the agencies from retaining control over their individual networks.

This is not to say that agencies may not choose to have their individual networks operated by the IRI, or even turned over to the IRI if they determine that to be appropriate.

Appropriate access control, privacy, and accounting mechanisms must be incorporated. This includes access control to data, resources, and the networks themselves, privacy of user data, and accounting mechanisms to support both cost allocation and cost auditing [23].

The technical and administrative approach must allow (indeed encourage) the incorporation of evolving technologies. In particular, the network must evolve towards provision of high bandwidth, type of service routing, and other advanced techniques to allow effective use of new computing technology in a distributed research environment.

Communications Infrastructure

The communications infrastructure provides connectivity between user machines, workstations, and centralized resources such as supercomputers and database machines. This roughly corresponds to communications services at and below the transport layer in the ISO OSI reference model. There are two different types of networks. The first are local networks, meaning those which are internal to a facility, campus, etc. The second are networks which provide transit service between facilities. These transit networks can connect directly to computers, but are evolving in a direction of connecting local networks. The networks supported by the individual agencies directly are mainly in the category of transit (or long-haul) networks, as they typically provide nationwide connectivity, and usually leave communications within a facility to be dealt with by the facility itself. The IRI communications infrastructure thus deals mainly with the interconnection of transit networks.

The internet model described above provides a simple method for interconnecting transit networks (as well as local networks.) By using IP gateways between the agency networks, packet transport service can be provided between computers on any of the various networks. The placement of the gateways and their capacity will have to be determined by an initial engineering study. In addition, as the IRI evolves, it may be cost-effective to install one or more wide area networks (or designate certain existing ones) to be IRI transit networks, to be used by all agencies on a cost sharing basis. Thus, the IRI communications infrastructure would consist of the interconnecting gateways plus any networks used specifically as transit networks. Using IP as the standard for interconnection of networks and global addressing provides a common virtual network packet transport service, upon which can be built various other network services such as file transfer and electronic mail. This will allow sharing of the communication facilities (channels, satellites, etc.) between the various user/agency communities in a cost effective manner.

To assure widespread interconnectivity, it is important that standards be adopted for use in the IRI and the various computers connected to it. These standards need to cover not only the packet transport capability but must address all the services required for networking in a scientific domain, including but not limited to file transfer, remote login, and electronic mail. Ultimately it is desirable to move towards a single set of standards for the various common services, and the logical choice for those standards are those being developed in the international commercial community (i.e. the ISO standards). However, many of the scientific networks today use one or more of a small number of different standards; in

particular the TCP/IP protocol suite mentioned above, the MFEnet protocols, and DECNET. As the international standards mature, it is expected that the number of communities using the same protocol suite will grow [5] [6]. Even today, several of the agencies/communities are using a common protocol suite, namely the TCP/IP suite. All the users connected to those computers and networks are able to have the full functions of an interoperable networking capability. And therefore the ability of the users to share resources and results will increase.

User Services

In order that scientists can effectively use the network, there needs to be a user support organization. To maximize the cost effectiveness of the overall IRI, the local user support personnel must be used effectively. In particular, it is anticipated that direct support of users/researchers would be provided by local support personnel. The IRI user support organization would provide support to those local support personnel in areas where nationwide common service is cost effective.

In particular, the this organization has several functions: assist the local support personnel in the installation of facilities compatible with the IRI, provide references to standard facilities (e.g. networking interfaces, mail software) to the local support personnel, answer questions that local personnel are not able to answer, aid in the provision of specific user community services, e.g. database of relevance to specific scientific domain.

Internet Research Coordination

To evolve internet to satisfy new scientific requirements and make use of new technology, research is required in several areas. These include high speed networking, type of service routing, new end to end protocols, and congestion control. The IRI organizational structure can assist in identifying areas of research where the various agencies have a common interest in supporting in order to evolve the network, and then assist in the coordination of that research.

MANAGEMENT APPROACH

A management approach is required that will allow each agency to retain control of its own networking assets while sharing certain resources with users sponsored by other agencies. To accomplish this, the following principles and constraints need to be followed.

IRI consists of the infrastructure to connect agency networks and

the user services required for effective use of the combined networks and resources.

An organization must be identified to be responsible for the engineering, operation, and maintenance of both the interconnecting infrastructure and the user services support.

While some agencies may choose to make use of IRI facilities and contractors to manage their individual agency networks, this would not be required and is not anticipated to be the normal situation. Any such arrangement would have to be negotiated individually and directly between the agency and the IRI operations organization. Normally, the IRI organization would neither manage the individual agency networks nor have any jurisdiction within such networks.

Gateways that interconnect the agency networks as well as any long-haul networks put in place specifically as jointly supported transit networks (if any such networks are required) will be managed and operated under the IRI organization.

A support organization for common IRI services is required. The principal clients for these services would be the local support personnel.

The IRI structure should support the coordination of the individual research activities required for evolution and enhancement of the IRI.

General Management Structure

Figure 2 shows the basic management structure for the IRI. It is based on the use of a non-profit organization (call it the Interagency Research Internet Organization, IRIO) to manage both the communications infrastructure and user support. The IRIO contracts for the engineering, development, operations, and maintenance of those services with various commercial and other organizations. It would be responsible for providing technical and administrative management of the contractors providing these functions. Having the IRI operational management provided by an independent non-profit organization skilled in the area of computer networking will permit the flexibility required to deal with the evolving and changing demands of scientific networking in a cost-effective manner.

Direction and guidance for the IRIO will be provided by a Policy Board consisting of representatives from the Government agencies who are funding the IRI. The Chairman of the Board will be selected from the agency representatives on a rotating basis. The Board will also have an Executive Director to provide administrative and other

support. To provide effective support for the IRI Policy Board as well as assure appropriate coordination with the IRIO, the Executive Director shall be the Director of the IRIO.

To assure that the IRI provides the best support possible to the scientific research community, the Policy Board will be advised by a Technical Advisory Board (TAB) consisting of representatives from the network research and engineering community, the various networks being interconnected with the IRI, and the scientific user community. Members of the TAB will be selected by the Policy Board. The TAB will review the operational support of science being provided by the IRI and suggest directions for improvement. The TAB will interface directly with the IRIO to review the operational status and plans for the future, and recommend to the Policy Board any changes in priorities or directions.

Research activities related to the use and evolution of the internet system will be coordinated by the Internet Research Activities Board (IRAB). The IRAB consists of the chairmen of the research task forces (see below) and has as ex-officio members technical representatives from the funding agencies and the IRIO. The charter of the IRAB is to identify required directions for research to improve the IRI, and recommend such directions to the funding agencies. In addition, the IRAB will continually review ongoing research activities and identify how they can be exploited to improve the IRI.

The Research Task Forces will each be concerned with a particular area/emphasis of research (e.g. end-to-end protocols, gateway architectures, etc.). Members will be active researchers in the field and the chairman an expert in the area with a broad understanding of research both in that area and the general internet (and its use for scientific research). The chairmen of the task forces will be selected by IRAB, and thus the IRAB will be a self-elected and governing organization representing the networking research community. The chairmen will solicit the members of the task force as volunteers.

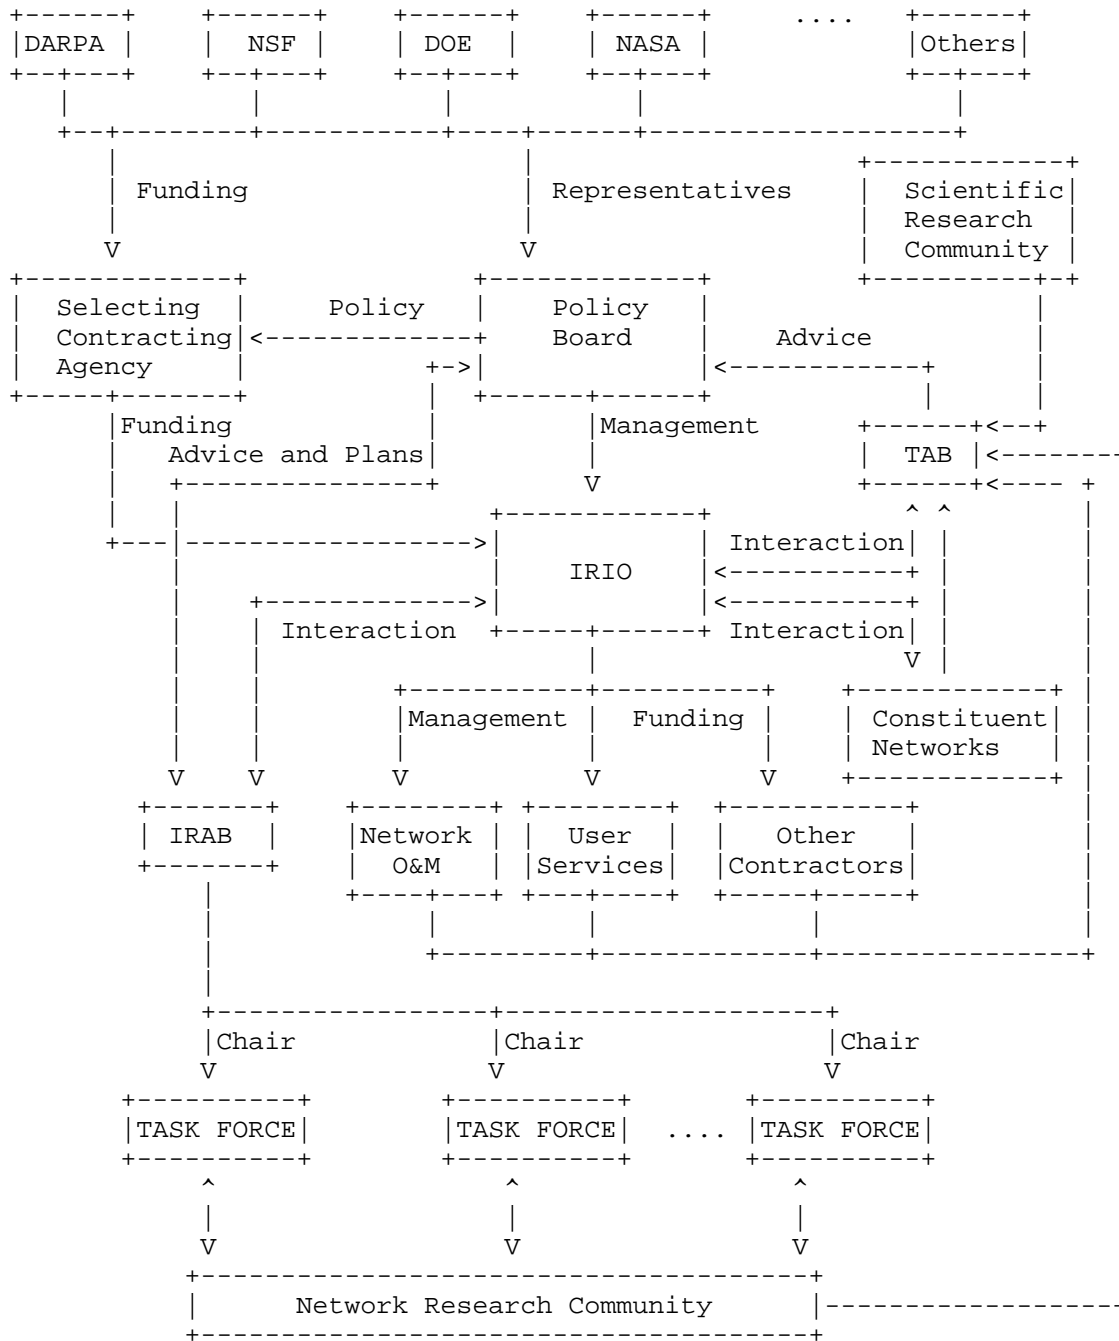


Figure 2: IRI Management Structure

Funding

In this section, the funding of the IRI is described. Recall that the IRI consists of the infrastructure to connect the agency networks and the services required for users to make effective use of such an infrastructure. These costs are divided into two categories; operations costs and research costs. The operations costs are those to operate and maintain both the communications infrastructure and the user services. These costs must be shared between the various agencies and channeled to the IRIO to operate the IRI. The research costs are those used to carry out the needed research to evolve the IRI. These costs are handled within the various agency budgets and used to support research in each agency with coordination between the agencies.

Operations Cost

Each participating agency will contribute a share of operations cost of IRI. Initially, each agency will contribute an equal share. Later, perhaps, the agency contributions will be adjusted according to a number of factors such as number of users, amount of traffic, type of support required (high bandwidth real time versus low bandwidth mail for example).

To facilitate the funding and administration of the IRI, one agency will be selected to manage the contract with IRIO. All funds will flow through that agency to the IRIO via interagency transfer. The role of the selected agency would be to provide the needed contractual activities and administrative management. Technical guidance and monitoring of IRIO activities would be provided by the IRI Policy Board.

It is not yet clear which Federal agency is best for this role. The requirements for such an agency include the ability to deal flexibly with the evolving requirements of the IRI, to deal with funding flowing from the various agencies, and to deal flexibly with the various agency technical representatives and incorporate their recommendations into the contract as required. One of the first activities required for the Policy Board would be to select an appropriate funding agency.

All operations and maintenance funding for the IRI will flow through the IRIO to selected contractors. This allows centralized management of the operation of the IRI.

There are two major assumptions underlying the budgetary estimates to follow. First of all, the IRIO should maintain a fairly low profile with respect to the end users (i.e. the scientists and

researchers). That is, the users will interact directly with their local support personnel. The IRIO will act as facilitator and coordinator, and provide facilities, information and help services to the local sites. This will allow the IRIO to remain relatively small, as it will not need to deal directly with the thousands of scientists/users.

Second, it is assumed that the operations budget supports the interconnection of agency networks as well as transit networking where required, but does not include costs of the individual agency networks.

Appendix A provides details of the budgetary estimate. Table 1 gives a summary. Note that the initial year has a higher expenditure of capital equipment, reflecting the need to purchase both the gateways needed for initial interconnection and the needed facilities to provide the operation of the gateways and the user services. Operations costs are expected to grow by inflation while the capital costs should remain constant (decrease when inflation is considered) as the IRI is stabilized.

Research Costs

In addition to the costs of operating and maintaining the communications infrastructure and user services, funding must be allocated to support an ongoing program of research to improve and evolve the IRI.

While each agency funds its own research program, the intent is that the various programs are coordinated through the IRI Policy Board. Likewise, while it is not intended that funds shall be combined or joint funding of projects is required, such joint activity can be done on an individual arrangement basis.

Each agency agrees, as part of the joint IRI activity, to fund an appropriate level of networking research in areas applicable to IRI evolution. The total funding required is currently estimated to be four million dollars in FY87, growing by inflation in the outyears. Details of this budgetary estimate are provided in Appendix A.

Table 1			
Annual IRI Operations Budget			
Fiscal Year	Capital Cost (\$M)	O & M Cost (\$M)	Total (\$M)
1987	2	8	10
1988	1	9	10
1989	1	10	11
1990	1	11	12
1991	1	12	13

PHASED IMPLEMENTATION PLAN

The long-term goal of the IRI activity is to put in place a functional high-performance network available to scientists across the nation. To accomplish this goal, a steady evolution of capability is envisioned. This phased approach involves both technical and administrative aspects.

Technical Phasing

Currently, networks are being supported by a number of agencies as discussed in Section 2. Many are using the DoD protocol suite (TCP/IP, etc.) and others have incorporated or are incorporating mechanisms for interoperability with networks using the DoD protocol suite (e.g. MFEnet). Most have discussed eventual evolution to ISO protocols and beyond. By and large, most of these networks are hooked together in some mainly ad hoc manner already, some by pairwise arrangement and some through third party connections (e.g. a university network connected to two agency networks).

There are two major shortcomings to this ad hoc connection, though. Performance is not adequate for advanced scientific environments, such as supercomputer usage, and community wide user support is not generally available. The phased approach described below will allow these deficiencies to be overcome through coordinated action on the part of the various funding agencies.

Phase I - Functional Interoperability

The initial stage of the IRI would provide for sharing of the communications facilities (e.g. channels, satellites, etc.) by interconnecting the networks using the Internet Protocol and IP gateways. In addition, mechanisms will be installed (where required) and maintained to allow interconnection of the common user services, such as electronic mail. This will allow sharing of resources attached to the network, such as supercomputers. [7] [8] Note: actual use of facilities other than mail would require arrangements with the various responsible parties for each host. For example, to login to a host not only requires network access; it also requires a login account on that host.

Specific steps to be undertaken in Phase I are the following:

Gateways will be purchased and installed where needed to interconnect the agency networks. The location and performance of these gateways will be specified by the IRIO and approved by the Policy Board. This engineering will take into account an estimate of current and future traffic requirements as well as existing interconnecting gateways. It may also result in a recommendation that some or all existing gateways between agency networks be replaced with common hardware so that adequate management of the interconnection can be achieved.

An IRI operations and management center will be established for the interconnecting gateways. [9] [10] This perhaps could be done in conjunction with a network management center for another set of gateways, e.g. those supported by DARPA or NSF.

The requirement for application gateways or other techniques to interconnect communities using different protocols will be investigated and a recommendation made by the IRIO in conjunction with the IRAB. The appropriate mechanisms will be installed by the IRIO at the direction of the Policy Board.

An initial user services facility will be established. This facility will provide at a minimum such services as a white pages of users (similar to the current Internet "whois" service) and a means for making accessible standard networking software.

The IRAB, in coordination with the Policy Board, will draft a coordinated research plan for the development of the new technologies required for evolution of the IRI.

Phase II - Full IRI Capability

Phase II will make the IRI fully functional with enhanced capabilities and performance.

High performance gateways with appropriate new capabilities and functions will be installed, replacing and/or augmenting the gateways in place from Phase I. The functionality and performance of these gateways will be specified based on the experience from Phase I use, the anticipated new uses of the network, and the state of the art technologies available as a result of the ongoing research.

The basic user services facility will be mature and support network operation. New capabilities will be developed to support specific scientific communities (such as a data base of software used by a specific community and its availability over the network.)

A high performance backbone network will be installed if needed to connect high performance agency networks. [11] [12] This is anticipated because of the move in several agencies to provide high bandwidth networks in support of such activities as supercomputer access.

The introduction and use of international standards will be investigated and a plan developed for providing more services to the broad scientific community through use of these standards.

Administrative Phasing

The goal of the IRI is to get to a fully cooperating and managed interagency research internet involving most if not all of the agencies supporting scientific research. Recognizing that currently, the major research networking players (both networking for research and research in networking) are DOE, NASA, DARPA, and NSF, the following steps are recommended:

The first and critical step is to establish a four agency Memorandum of Agreement (MOA) to interconnect the agency networks and to share the costs of interconnection, transit networks, and an operations center. A management structure should be agreed upon as outlined above. Agreement must also be reached on the need to fund an ongoing research and engineering activity to evolve the internet.

A Policy Board and Technical Advisory Board should be established as quickly as possible to assure appropriate guidance and direction.

The Policy Board shall then select an agency to handle the

administrative and contractual actions with the IRIO.

A non-profit organization shall then be selected by that agency through an appropriate procurement mechanism to be the IRIO. The Policy Board of the IRI shall be the selection panel.

The initial four agencies shall transfer the agreed upon funds to the selected contracting agency on equal basis to start.

These funds will then allow the contracting agency to establish a contract for the IRIO with the selected non-profit organization.

The IRIO can then establish sub-contracts for engineering, procurement, installation, and management of gateways and operation of the user services center.

To initiate the research coordination, the following steps will be accomplished.

The Internet Activities Board will evolve into the Internet Research Activities Board, through added membership and charter revision.

Additional task forces will be formed as needed to reflect the expanded areas of research interest.

Once the IRI is established and operating, the funding and use of the IRI will be reviewed to determine if equal funding is equitable. If not, the IRIO should be tasked to develop a recommendation for a practical cost allocation scheme. In addition, once the IRI has proved itself to be successful, other agencies will join the IRI and provide additional funding.

INDUSTRY ROLE

This report has thus far addressed the interconnection of agency supported networks and the use of such an internet by agency supported researchers. However, industry also has a need for a similar infrastructure to support its research activities. [13] [14]. Note that this refers only to industrial research activities. It is not envisioned, nor would it be appropriate, for the IRI to provide a communications system for normal industrial activities. Regulatory concerns make it difficult for industry to connect to a network that is supported by a federal agency in pursuit of the agency mission.

The IRI structure above, though, may permit the connection of industrial research organizations. Since the IRIO is a non-profit non-government organization, it would be able to accept funds from

industry as a fair share of the costs of using the IRI. These funds in turn can be used to expand the networking resources so that no degradation of service is felt by the users supported by the federal agencies. This topic would need to be discussed further by the Policy Board and the organization selected as the IRIO.

SUMMARY AND CONCLUSIONS

The interconnection of the various agency networks supporting scientific research into an overall infrastructure in support of such research represents an exciting opportunity. This report recommends an approach and a specific set of actions that can achieve that goal. It is hoped that, regardless of the mechanism used, that the Federal agencies involved recognize the importance of providing an appropriate national infrastructure in support of scientific research and take action to make such an infrastructure a reality.

ACKNOWLEDGEMENT

This report was prepared with advice and comments from a large number of people, including the members of the FCCSET Committee Network Working Group and the Internet Activities Board. Their input is greatly appreciated, and I hope that this report represents a consensus on both the need for the IRI and the proposed approach.

APPENDIX A - FUNDING BREAKDOWN

This appendix provides the details for the budgetary estimates of Table 1.

Gateways

Gateways will be required between the various agency (and perhaps regional) networks. As an upper bound, assume one IRI gateway per state times \$40K per gateway, spread out over two years, for a capital cost of \$1M per year for first two years.

Operation Center

The IRI operations center will have to engineer the location and capacity of the gateways, as well as install, operate and maintain them. It also will need to coordinate support and maintenance of end-to-end service, helping to identify and correct problems in the interconnections. Costs are estimated as two people round the clock to man the operations center and three full time people to coordinate, operate, and engineer the IRI. Using an estimate of \$120K (including other direct costs (ODC)) per year for an operator and \$200K per year for other activities, and translating 2 people round the clock into 9 people results in a total annual cost of \$1.7M. In addition, equipment costs of roughly \$500K per year can be expected.

Transit Networks

It is expected that support of at least one transit network will be necessary. This may involve reimbursement to one of the agencies for use of their network, or may involve operations and maintenance of an IRI dedicated network. An estimate for these costs, based on historical data for operating the Arpanet, is \$4M per year.

User Support Organization

To provide effective support as discussed above will require a staff available during working hours. A reasonable estimate for the costs of such an organization is 5 people times \$200K per year, or \$1M per year (including ODC). In addition, there will be capital equipment costs in the first two years totalling roughly \$2M.

REFERENCES

1. FCCSET Committee on Very High Performance Computing Network Working Group, Report on Interagency Networking for Research Programs, February 1986.
2. Cerf, V.G. and P. Kirstein, "Issues in packet-network interconnection," Proceedings of the IEEE, pp. 1386-1408, November 1978
3. Cerf, V.G. and E. Cain, "The Dod internet architecture model," Computer Networks, pp. 307-318, July 1983.
4. Leiner, B.M., J. Postel, R. Cole, and D. Mills, "The DARPA internet protocol suite," IEEE communications Magazine March 1985.
5. Defense Advanced Research Projects Agency, A History of the Arpanet: The First Decade, Defense Advanced Research Projects Agency, April 1981. (Defense Tech. Info. Center AD A1 15440)
6. Jacobs, I.M. et. al., "General purpose satellite networks," Proceedings of the IEEE pp. 1448-1467, November 1978
7. Tobagi, F., R. Binder, and B.M. Leiner, "Packet radio and satellite networks," IEEE Communications Magazine, November 1984.
8. Kahn, R.E. et. al., "Advances in packet radio technology," Proceedings of the IEEE pp. 1468-1496, November
9. Clark, D. et. al., "An introduction to local area networks," Proceedings of the IEEE, November 1978
10. Lederberg, J., "Digital communications and the conduct of science: the new literacy," vol. 66, pp. 1314-1319, November 1978.
11. Hoskins, J.C. and J.S. Quaterman, "Notable Computer Networks," pp. 932-971, October 1986.
12. Dennings, P.J., A.C. Hearn, and C.W. Kern, "History and overview of CSNET," pp. 138-145, March 1983.
13. Comer, D., "The computer science research network CSNET: A history and status report", vol. 26, pp. 747-753, October 1983.

14. Bailey, R.R. NAS: supercomputing master tool for aeronautics Aerospace America, pp. 118-121, January 1985
15. Jennings, D.M., L.H. Landweber, I.H. Fuchs, W.R. Adrion "Computer Networking for Scientist Science" vol. 231 pp. 943-950, February 1986
16. Cerf, V.G. R.E. Kahn, "A protocol for packet network intercommunication, IEEE Transactions on Communications vol. COM-22, May 1974
17. Zimmerman, H. "OSI reference model - the ISO model of architecture for open systems intercommunications, IEEE Transactions on Communications vol. COM-28 pp. 425-432 April 1980
18. Defense Communications Agency, MIL STD 1777: Internet Protocol, 1983
19. Defense Communications Agency, MIL STD 1778: Transmission Control Protocol Defense Communications Agency, 1983
20. Defense Communications Agency, MIL STD 1780: File Transfer Protocol Defense Communications Agency, 1985
21. Defense Communications Agency, MIL STD 1781: Simple Mail Transfer Protocol Defense Communications Agency, 1985
22. Defense Communications Agency, MIL STD 1782: Telnet Protocol Defense Communications Agency, 1985
23. Leiner, B.M. and M. Bishop, Research Institute for Advanced Computer Science Access Control and Privacy in Large Distribution Systems, RIACS TR 86.6, March 1986