

THOUGHTS ON THE NATIONAL RESEARCH AND EDUCATION NETWORK

Status of this Memo

The memo provides a brief outline of a National Research and Education Network (NREN). This memo provides information for the Internet community. It does not specify any standard. It is not a statement of IAB policy or recommendations.

Distribution of this memo is unlimited.

ABSTRACT

This contribution seeks to outline and call attention to some of the major factors which will influence the form and structure of a National Research and Education Network (NREN). It is implicitly assumed that the system will emerge from the existing Internet.

ACKNOWLEDGEMENTS

The author gratefully acknowledges support from the National Science Foundation, The Defense Advanced Research Projects Agency, the Department of Energy and the National Aeronautics and Space Administration through cooperative agreement NCR-8820945. The author also acknowledges helpful comments from colleagues Ira Richer, Barry Leiner, Hans-Werner Braun and Robert Kahn. The opinions expressed in this paper are the personal opinions of the author and do not represent positions of the U.S. Government, the Corporation for National Research Initiatives or of the Internet Activities Board. In fact, the author isn't sure he agrees with everything in the paper, either!

A WORD ON TERMINOLOGY

The expression "national research and education network" is taken to mean "the U.S. National Research and Education Network" in the material which follows. It is implicitly assumed that similar initiatives may arise in other countries and that a kind of Global Research and Education Network may arise out of the existing international Internet system. However, the primary focus of this paper is on developments in the U.S.

FUNDAMENTALS

1. The NREN in the U.S. will evolve from the existing Internet base. By implication, the U.S. NREN will have to fit into an international environment consisting of a good many networks sponsored or owned and operated by non-U.S. organizations around the world.
2. There will continue to be special-purpose and mission-oriented networks sponsored by the U.S. Government which will need to link with, if not directly support, the NREN.
3. The basic technical networking architecture of the system will include local area networks, metropolitan, regional and wide-area networks. Some nets will be organized to support transit traffic and others will be strictly parasitic.
4. Looking towards the end of the decade, some of the networks may be mobile (digital, cellular). A variety of technologies may be used, including, but not limited to, high speed Fiber Data Distribution Interface (FDDI) nets, Distributed-Queue Dual Bus (DQDB) nets, Broadband Integrated Services Digital Networks (B-ISDN) utilizing Asynchronous Transfer Mode (ATM) switching fabrics as well as conventional Token Ring, Ethernet and other IEEE 802.X technology. Narrowband ISDN and X.25 packet switching technology network services are also likely play a role along with Switched Multi-megabit Data Service (SMDS) provided by telecommunications carriers. It also would be fair to ask what role FTS-2000 might play in the system, at least in support of government access to the NREN, and possibly in support of national agency network facilities.
5. The protocol architecture of the system will continue to exhibit a layered structure although the layering may vary from the present-day Internet and planned Open Systems Interconnection structures in some respects.
6. The system will include servers of varying kinds required to support the general operation of the system (for example, network management facilities, name servers of various types, email, database and other kinds of information servers, multicast routers, cryptographic certificate servers) and collaboration support tools including video/teleconferencing systems and other "groupware" facilities. Accounting and access control mechanisms will be required.
7. The system will support multiple protocols on an end to end basis. At the least, full TCP/IP and OSI protocol stacks will be supported. Dealing with Connectionless and Connection-Oriented Network Services in the OSI area is an open issue (transport service bridges and

application level gateways are two possibilities).

8. Provision must be made for experimental research in networking to support the continued technical evolution of the system. The NREN can no more be a static, rigid system than the Internet has been since its inception. Interconnection of experimental facilities with the operational NREN must be supported.

9. The architecture must accommodate the use of commercial services, private and Government-sponsored networks in the NREN system.

Apart from the considerations listed above, it is also helpful to consider the constituencies and stakeholders who have a role to play in the use of, provision of and evolution of NREN services. Their interests will affect the architecture of the NREN and the course of its creation and evolution.

NREN CONSTITUENTS

The Users

Extrapolating from the present Internet, the users of the system will be diverse. By legislative intent, it will include colleges and universities, government research organizations (e.g., research laboratories of the Departments of Defense, Energy, Health and Human Services, National Aeronautics and Space Administration), non-profit and for-profit research and development organizations, federally funded research and development centers (FFRDCs), R&D activities of private enterprise, library facilities of all kinds, and primary and secondary schools. The system is not intended to be discipline-specific.

It is critical to recognize that even in the present Internet, it has been possible to accommodate a remarkable amalgam of private enterprise, academic institutions, government and military facilities. Indeed, the very ability to accept such a diverse constituency turns on the increasing freedom of the so-called intermediate-level networks to accept an unrestricted set of users. The growth in the size and diversity of Internet users, if it can be said to have been constrained at all, has been limited in part by usage constraints placed on the federally-sponsored national agency networks (e.g., NSFNET, NASA Science Internet, Energy Sciences Net, High Energy Physics Net, the recently deceased ARPANET, Defense Research Internet, etc.). Given the purposes of these networks and the fiduciary responsibilities of the agencies that have created them, such usage constraints seem highly appropriate. It may be beneficial to search for less

constraining architectural paradigms, perhaps through the use of backbone facilities which are not federally-sponsored.

The Internet does not quite serve the public in the same sense that the telephone network(s) do (i.e., the Internet is not a common carrier), although the linkages between the Internet and public electronic mail systems, private bulletin board systems such as FIDONET and commercial network services such as UUNET, ALTERNET and PSI, for example, make the system extremely accessible to a very wide variety of users.

It will be important to keep in mind that, over time, an increasing number of institutional users will support local area networks and will want to gain access to NREN by that means. Individual use will continue to rely on dial-up access and, as it is deployed, narrow-band ISDN. Eventually, metropolitan area networks and broadband ISDN facilities may be used to support access to NREN. Cellular radio or other mobile communication technologies may also become increasingly popular as access tools.

The Service Providers

In its earliest stages, the Internet consisted solely of government-sponsored networks such as the Defense Department's ARPANET, Packet Radio Networks and Packet Satellite Networks. With the introduction of Xerox PARC's Ethernet, however, things began to change and privately owned and operated networks became an integral part of the Internet architecture.

For a time, there was a mixture of government-sponsored backbone facilities and private local area networks. With the introduction of the National Science Foundation NSFNET, however, the architecture changed again to include intermediate-level networks consisting of collections of commercially-produced routers and trunk or access lines which connected local area network facilities to the government-sponsored backbones. The government-sponsored supercomputer centers (such as the National Aerospace Simulator at NASA/AMES, the Magnetic Fusion Energy Computing Center at Lawrence Livermore Laboratory and the half-dozen or so NSF-sponsored supercomputer centers) fostered the growth of communications networks specifically to support supercomputer access although, over time, these have tended to look more and more like general-purpose intermediate-level networks.

Many, but not all, of the intermediate-level networks applied for and received seed funding from the National Science Foundation. It was and continues to be NSF's position, however, that such

direct subsidies should diminish over time and that the intermediate networks should become self-sustaining. To accomplish this objective, the intermediate-level networks have been turning to an increasingly diverse user constituency (see section above).

The basic model of government backbones, consortium intermediate level nets and private local area networks has served reasonably well during the 1980's but it would appear that newer telecommunications technologies may suggest another potential paradigm. As the NSFNET moves towards higher speed backbone operation in the 45 Mb/s range, the importance of carrier participation in the enterprise has increased. The provision of backbone capacity at attractive rates by the inter-exchange carrier (in this case, MCI Communications Corporation) has been crucial to the feasibility of deploying such a high speed system.

As the third phase of the NREN effort gets underway, it is becoming increasingly apparent that the "federally-funded backbone" model may and perhaps even should or must give way to a vision of commercially operated, gigabit speed systems to which the users of the NREN have access. If there is federal subsidy in the new paradigm, it might come through direct provision of support for networking at the level of individual research grant or possibly through a system of institutional vouchers permitting and perhaps even mandating institution-wide network planning and provision. This differs from the present model in which the backbone networks are essentially federally owned and operated or enjoy significant, direct federal support to the provider of the service.

The importance of such a shift in service provision philosophy cannot be over-emphasized. In the long run, it eliminates unnecessary restrictions on the use and application of the backbone facilities, opening up possibilities for true ubiquity of access and use without the need for federal control, except to the extent that any such services are considered in need of regulation, perhaps. The same arguments might be made for the intermediate level systems (metropolitan and regional area access networks). This does NOT mean that private networks ranging from local consortia to inter-continental systems will be ruled out. The economics of private networking may still be favorable for sufficiently heavy usage. It does suggest, however, that achieving scale and ubiquity may largely rely on publicly accessible facilities.

The Vendors

Apart from service provision, the technology available to the users and the service providers will come largely from commercial sources. A possible exception to this may be the switches used in the gigabit testbed effort, but ultimately, even this technology will have to be provided commercially if the system is to achieve the scale necessary to serve as the backbone of the NREN.

An important consequence of this observation is that the NREN architecture should be fashioned in such a way that it can be constructed from technology compatible with carrier plans and available from commercial telecommunications equipment suppliers. Examples include the use of SONET (Synchronous Optical Network) optical transmission technology, Switched Multimegabit Data Services offerings (metropolitan area networks), Asynchronous Transmission Mode (ATM) switches, frame relays, high speed, multi-protocol routers, and so on. It is somewhat unclear what role the public X.25 networks will play, especially where narrow and broadband ISDN services are available, but it is also not obvious that they ought to be written off at this point. Where there is still research and development activity (such as in network management), the network R&D community can contribute through experimental efforts and through participation in standards-making activities (e.g., ANSI, NIST, IAB/IETF, Open NMF).

OPERATIONS

It seems clear that the current Internet and the anticipated NREN will have to function in a highly distributed fashion. Given the diversity of service providers and the richness of the constituent networks (as to technology and ownership), there will have to be a good deal of collaboration and cooperation to make the system work. One can see the necessity for this, based on the existing voice network in the U.S. with its local and inter-exchange carrier (IEC) structure. It should be noted that in the presence of the local and IEC structure, it has proven possible to support private and virtual private networking as well. The same needs to be true of the NREN.

A critical element of any commercial service is accounting and billing. It must be possible to identify users (billable parties, anyway) and to compute usage charges. This is not to say that the NREN component networks must necessarily bill on the basis of usage. It may prove preferable to have fixed access charges which might be modulated by access data rate, as some of the intermediate-level networks have found. It would not be surprising to find a mixture of charging policies in which usage charges are preferable for small

amounts of use and flat rate charges are preferred for high volume use.

It will be critical to establish a forum in which operational matters can be debated and methods established to allow cooperative operation of the entire system. A number of possibilities present themselves: use of the Internet Engineering Task Force as a basis, use of existing telecommunication carrier organizations, or possibly a consortium of all service providers (and private network operators?). Even if such an activity is initiated through federal action, it may be helpful, in the long run, if it eventually embraces a much wider community.

Agreements are needed on the technical foundations for network monitoring and management, for internetwork accounting and exchange payments, for problem identification, tracking, escalation and resolution. A framework is needed for the support of users of the aggregate NREN. This suggests cooperative agreements among network information centers, user service and support organizations to begin with. Eventually, the cost of such operations will have to be incorporated into the general cost of service provision. The federal role, even if it acts as catalyst in the initial stages, may ultimately focus on the direct support of the users of the system which it finds it appropriate to support and subsidize (e.g., the research and educational users of the NREN).

A voucher system has been proposed, in the case of the NREN, which would permit users to choose which NREN service provider(s) to engage. The vouchers might be redeemed by the service providers in the same sort of way that food stamps are redeemed by supermarkets. Over time, the cost of the vouchers could change so that an initial high subsidy from the federal government would diminish until the utility of the vouchers vanished and decisions would be made to purchase telecommunications services on a pure cost/benefit basis.

IMPORTANCE OF COMMERCIAL INTERESTS

The initial technical architecture should incorporate commercial service provision where possible so as to avoid the creation of a system which is solely reliant on the federal government for its support and operation. It is anticipated that a hybrid system will develop but, for example, it is possible that the gigabit backbone components of the system might be strictly commercial from the start, even if the lower speed components of the NREN vary from private, to public to federally subsidized or owned and operated.

CONCLUSIONS

The idea of creating a National Research and Education Network has captured the attention and enthusiasm of an extraordinarily broad collection of interested parties. I believe this is in part a consequence of the remarkable range of new services and facilities which could be provided once the network infrastructure is in place. If the technology of the NREN is commercially viable, one can readily imagine that an economic engine of considerable proportions might result from the widespread accessibility of NREN-like facilities to business sector.

Security Considerations

Security issues are not discussed in this memo.

Author's Address

Vinton G. Cerf
Corporation for National Research Initiatives
1895 Preston White Drive, Suite 100
Reston, VA 22091

EMail: vcertif@NRI.Reston.VA.US

Phone: (703) 620-8990