The COOK Report on Internet → NREN

Building New Backbones : A Tale of Two Philosophies

Vol. I, Nos. 6 & 7, Sept-Oct 1992

NSF Appears Determined to Repeat Past Errors While NASA & DOE Execute Outstanding Backbone Rebid

A Funny Thing Happened on the Way to the NREN

NSF No Longer Lead Agency

We used to think that the NSF and its hand-picked partner ANS were to be responsible for the design and creation of the NREN. After all, S272 passed in on September 11 1991 with the statement that the National Science Foundation shall "act as lead agency in coordinating the collaboration among federal agencies contributing to the network." What was less noticed was the language in HR 656 where the Director of the Office of Science and Technology Policy was to "coordinate the implementation of agency and departmental activities supporting the broad deployment and use of a national multi-gigabit per second research and education computer network. . ." The final language of the HPCC Act PL 102 -194 states "the National Science Foundation, the Department of Defense, the Department of Energy, the Department of Commerce, the National Aeronautics and Space Administration and other agencies participating in the program shall support the establishment of the National Research and Education Network." The Director of OSTP is no longer explicitly listed as Coordinator for network development. Instead the law states that the President shall implement "a National High Performance Computing Program which shall (a) establish the goals and priorities for federal high performance computing research, development, networking and other activities and (b) provide for interagency coordination of ... networking. (2) The program shall ... provide for the oversight of the operation and evolution of the Network. . . . ." The origin of the Program is from the President of the United States. It is presumably created and provided for his acceptance by the Director of OSTP. So again in a somewhat more roundabout way, the Director of OSTP is the focus for coordination and leadership within NREN not the NSF.

Now consider what Public Law 102-194 does tell the NSF to do. Sec. 201 (a) states: "The National Science Foundation shall provide computing and networking infrastructure support for all science and engineering disciplines and support basic research and human resource development in all aspects of high performance computing and advanced high speed computer networking."

Networking infrastructure support is basic, plain vanilla connectivity. Nothing here to do with pushing the gigabit state-of-the-art envelope. "Support basic research and human resource development in all aspects of high speed

(Cont. on page 2.)

Special Double Issue

Because of a trip to Russia in October we shall publish only two issues duringing the months of November, December and January.

IRS States That as of August 25, 1992 ANS Does not Have 501(c)(3) Status

In July I identified a source in the New York City Public Affairs Office of the Internal Revenue Service. I sent him a written request to be informed as to Advanced Network and Service's 501 (c)(3) status. On August 7th and again on August 25th he stated that ANS did not have 501(c)(3) status. All the IRS could do he said was tell me whether ANS had the status or not. It did not have it. This could mean, he said one of five possibilities:

1. They had never applied for it, but since they had told me that they did apply on October 27, 1991, this was unlikely.

(Cont. on page 19.)

On the Inside

computer networking" reads like: assist DARPA's network development efforts and anyone's else efforts to make the network easier to use. Nothing wrong with that either. What does not seem to be a part of the marching orders is "develop your own state-of-the-art technology-pushing backbone."

Sec. 201 (a) also states that, to the extent that colleges, universities and libraries are unable to connect to the network with private sector assistance, the National Science Foundation shall serve as a connector of last resort. The NSF is also charged with the responsibility of "serving as the primary source of information on access to and use of the network."

Finally the NSF is instructed to "upgrade the National Science Foundation funded network, assist regional networks to upgrade their capabilities, and provide other Federal departments and agencies the opportunity to connect to the network." "Upgrade the . . . network" is probably meant in the sense of make sure the network carrying capacity expands so the agency's basic mission of connecting researchers and educators is carried out. The meaning "upgrade" in such a general statement is certainly open to interpretation.

The Likely Roles

Recent events have shown the likely role of the National Science Foundation in NREN. There will be the NSFnet designed to ensure, by subsidization if necessary, broad connectivity. Milnet will be likely to continue to perform its current role. ESnet and NSI net will be the networks where the state-of-the-art envelope is most like to be pushed in other than proprietary ways. In the nine months since HPCC passed, the NSF's way of network building has gotten it firmly embroiled in controversy and has enabled it to do nothing definitive towards the carrying out of HPCC mandates. All we have is the solicitation draft for another cost sharing follow-on to the current NSFnet. The comments are predominantly negative. (One respondent actually suggested the NSF revise the draft as a draft and reissue it before going to a Final Solicitation!) In the meantime contracts for upgraded ESnets and NSInets are being let. Furthermore the Department of Energy has announced a series of innovative consortia agreements to attack practical grand challenge applications with massively parallel supercomputers. (See accompanying article: DOE Announces Innovative Grand Challenge Supercomputing Projects, page 8.)

Two Ways to Build Backbones

Two distinct and opposed philosophies of backbone building have emerged. "Cost share" in a cooperative agreement with few specifications most of which leave the awardee plenty of room for experimentation and an advantage for technology exploitation in the marketplace that the cost sharing partner "buys into." This is the NSF way of doing business.

The second philosophy is a direct bid process for a clearly defined technical contract where the selection panel makes it decision on technical excellence with bidder cost shielded out of the initial review process. The goal is to purchase sate-of-the-art but also COTS (Commercial off the Shelf) with the contract awardee not expected to use technology developed specifically for the project.

The second philosophy is a direct bid process for a clearly defined technical contract where the selection panel makes it decision on technical excellence with bidder cost shielded out of the initial review process. The goal is to purchase sate-of-the-art but also COTS (Commercial off the Shelf) with the contract awardee not expected to use technology developed specifically for the project. According to NASA's Milo Medin, a NSInet manager and engineer directly involved in the upgrade: "We gain the most leverage for our program dollars by working in directions that the carriers already are going in, but serving as an incentive for them to accelerate progress in these areas."

"Cost Sharing:" The NSFnet Way

The current five year cooperative agreement with MERIT is an example. When it "ends" in November, the NSF will have paid about 28 to 30 million dollars to MERIT and it's joint study partners IBM and MCI now operating as ANS. In return the NSF has thrown out estimates ranging from 30 to 100 million dollars as the amounts that the awardees have actually spent in providing the NSFnet service. The July and August issues of the *COOK Report* have already presented critiques of the NSF's new backbone rebid solicitation -- one that promises to repeat the errors of the earlier MERIT agreement. Some other interesting opinions have since become available.

The Electronic Frontier Foundation runs a periodic Communications Policy Forum. On July 23rd, one such forum "composed 25 representatives from industry, government and the non-profit sector met to discuss the elements of the Draft Solicitation and the issues it raised. . . . . The discussion, revealed generally the lack of shared understanding about the underlying elements of the Draft, i.e., what is a NAP, what is the vBNS, and how will they relate, and the importance of the many points at which NSF has remained 'intentionally silent' on crucial issues that will affect the deployment of the network." (If the NSF is truely interested in an open and fair solicitation process, one would wonder why it would remain "intentional silent" on any critical issues? -- Quoted passages describing meeting are from Memo of meeting sent by EFF to NSF as its critique of solicitation.)

"Steve Wolff of NSF provided an overview of the Draft Solicitation itself and answered questions from roundtable participants. He noted that the Draft Solicitation invited specific comment on architecture, but little in the way of its impact on stakeholders. He suggested that the draft's most "glaring omission" is that there are no rules by which the submission's will be evaluated, and stated that it would be useful to hear thoughts on how to evaluate submissions."

Solicitation Purpose Still not Clear

"Wolff described the NAPs as open to all traffic and as the playing field through which all may bring their networks. [Editor - NAPS have a requirement for the transmission of video so right away that eliminates some small providers.] The vBNS will operate at 155 mb/s, a figure chosen because applications are de-

manding it and because 155 may be the current limit of affordability. [Wolff apparently was not specific about what these applications are and as we shall see in a moment one of the NSF's critics believes that the 155 figure is unwarrentedly high -- chosen for PR purposes. - Editor] The only special role for the vBNS is that it will carry the R&E traffic that fits the NSF's acceptable use policies. However, anyone who wants to connect to the NAPs could provide separate high-speed 'backbone' service."

"The relationship between the NAPs and the vBNS [is unclear] -- Are NAPs the only connection point to the vBNS, or could other high-speed backbone service providers connect to the vBNS at the "most sensible place" for them? According to Wolff, NAPs will be the only interconnection point with the vBNS at the start, but the Draft seems to contemplate other connections in the future. [Such apparently is the squishy language of a cooperative agreement especially when the NSF (the King) is trying to sell Lettres Patent to the highest bidder for exploitation of the network - Editor.] NSF believes that there will be other providers of high speed networking, but it has no policy for how they will connect."

Apparently a lot of key elements still haven't been thought through. "In light of the prospect of a physically shared, logically partitioned high speed backbone, commercial providers asked whether the vBNS provider would be able to offer direct connections to the vBNS. NSF staff agreed that they did not know whether that would happen."

Apparently a lot of key elements still haven't been thought through. "In light of the prospect of a physically shared, logically partitioned high speed backbone, commercial providers asked whether the vBNS provider would be able to offer direct connections to the vBNS. NSF staff agreed that they did not know whether that would happen." "There is nothing in the draft that precludes the creation of non-NSF sponsored NAPs that create their own relationship with NSF-sponsored NAPs. The solicitation does not attempt to define the entire universe of network interconnection points." [This seems absurd. Under cost sharing the NSF chooses apparently to set almost no conditions on how the winner gets to exploit the 'cooperative agreement?' - editor]

"NAP location and number will affect connection costs and could create geographically based advantages for some service providers. [The concern was raised that] who is allowed to connect to the NAPs may affect where they are placed. Some participants asked NSF to clarify when such decisions would be made, [and] if not [to] specify how many NAPs should be provided or what criteria bidders should address in proposing NAP placement." Another concern was "Equity of access to networking bandwidth: If there are a small number of NAPs, then NAP location becomes critical because it will tend to favor some institutions over others. The result could be institutions disadvantaged because of their location."

[Other questions:] "Starting a private, non-NSF sponsored NAP and connecting to the vBNS -- how it would be done, what the connection charges would be, and how they would get determined? NSF staff has no official position on that issue, but suggested that if commenters believed that NSF ought to address the issue, they should say so in their response to the Draft Solicitation."

"A number of participants suggested that instead of an either/or proposition, the real issue comes in determining the handoff between the two, [commercial networks connect via the NAPs to the vBNS], as "precompetitive" or experimental elements migrate into the off-the-shelf production network. A request was made that NSF clarify related issues: how this transition will work, what determines when the NSF pulls out of an arrangement in order to free the private sector, and what vehicle will exist to transfer technology from experimental use to production markets."

"NAP attachment prices -- Wolff said he would like to see a pricing structure as part of a proposal. His guess is that some large but not\dominant fraction of the NAP/RA cost of doing business should be recoverable through the fee structure. [Meanwhile, other criticisms pointed out that] "the NSF should realize and acknowledge that: 'the proposed architecture does not resolve the AUP issue, but simply moves it.'"

Commercial Network Providers

"The perspective from commercial service providers ranged widely from general support to condemnation."

"Concerns:

> The solicitation contradicts the direction of the High Performance Computing Act.
>
> The connection subsidies essentially give some competitors an unfair advantage in an otherwise commercial marketplace.
>
> The CIX can already provide ubiquity, so that the argument that the NSFNet is needed for widespread connectivity is weak.
>
> The government's provision of high speed networks for R&E unfairly competes with commercial providers.
>
> If the vBNS is provided over shared facilities, there will not be a level playing field for vBNS access."

"More generally, it was argued [by the telcos attending] that the NSF is essentially prescribing a new commercial network architecture that will operate as a common carrier and recover its costs from its users. The NSF should not be prescribing network architecture for a commercial common carrier, since those matters are handled by the Federal Communications Commission."

"More generally, it was argued [by the telcos attending] that the NSF is essentially prescribing a new commercial network architecture that will operate as a common carrier and recover its costs from its users. The NSF should not be prescribing network architecture for a commercial common carrier, since those matters are handled by the Federal Communications Commission."

"One suggestion was made in an attempt to address these concerns: that a consortium of providers that included both IXCs and LECs jointly bid on the solicitation. Such a consortium, it was suggested, could allow LECs to

participate in local transport and switching while IXCs handle inter-LATA transport; create incentives for participants to address the competitive "level playing field" issues equitably, since they would be partners; and allow the risk for developing the high speed network to be broadly shared." [A good idea but people shouldn't hold their breath waiting for it to happen.]

FARnet Reactions Negative

A couple of days later FARnet (Federation of American Research Networks) had a meeting at which according to an article in *Network World* (July 27th), and a memo from a telecom executive, emphasized many additional points -- among them:

a request in the final solicitation for "the creation of an advisory committee to oversee the management of the new backbone net so that the NSF design concept can evolve gracefully into reality."

a concern about a "plan to require network providers attached to the NAPs to support video since no standard for video-conferencing using the internet protocol exists today."

a recommendation that the NSF require winning bidders to provide "a transition plan from the current T-3 network to the new routing scheme for the sake of regional service providers."

Network World (7/27) also noted that a service provider from an RBOC stated that "by setting up the NAPs, the federal government begins to play a role in defining the infrastructure of the newly emerging market for commercial IP data services."

"by setting up the NAPs, the federal government begins to play a role in defining the infrastructure of the newly emerging market for commercial IP data services."

According to the memo received from a telecom executive FARnet also wanted "greater emphasis on ubiquity and less on high-end networks." In other words FARnet was saying make sure that you don't design something so fancy and that our less well-to-do customers can't afford to pay for it. Also the NSF should take great care with the location of NAPs which if there weren't enough could affect both academic institutions and telcos negatively.

EDUCOM's NTTF Comments

Educom's National Telecommunications Task Force, which has been very sympathetic to MERIT and the NSF, chimed in several comments. Here are a few:

"The draft solicitation does not adequately provide a context for the next iteration of NSFNET and its relationship to other federal programs such as the IINREN, Gigabit testbed program, NREN itself, commercialization, and the rapidly developing call for a national information infrastructure. Since the funds proposed to be allocated to

For the Layman -- A Quick Tutorial on the Difference Between a Switch and a Router

I asked Milo Medin to comment on this he came back with the following reply.

"The ATM switch does not operate at the IP level, it operates at a lower level, like Ethernet or FDDI (Fiber distributed Data Interface 100M bit per second LAN). If you have traditional hosts on something like an FDDI, they still need to go through a router to talk through an ATM switch. So, in our plans, we have Cisco IP routers which attach to FDDI rings, and then there is a DS-3 (initially, SONET later) that also plugs into the router and connects at the other end to an ATM switch. [This is how we "connect" the TCP/IP and application layers to level one and two (ATM).]

The switch is a switch because it can take data from say port A and switch it in hardware to Port C, and from port B to Port D all at the same time. [The hardware that does this is called the switching "fabric".] Conventional routers and bridges take a packet from Port A, and move the data across a bus or backplane to another port. The switch doesn't have a central backplane or bus, it has [a very large number of] parallel paths so all those ATM cells can be switched in parallel from port to port. That lets you build switches with huge number of ports and capacity, since there isn't a central [backplane or bus] bottleneck. Even in a Cisco, that central bus operates at 522 Mbps. That's fine for a few interfaces but a telco needs to support hundreds.

Another difference is that the switch routes in hardware. That is, the ATM cell header is very simple, and all the switch decides is which port to send a cell out on, all in hardware. There is no way to really support the cell switching rates needed for 155 and 622 without putting a lot of this into hardware. ATM cells are designed to support this kind of hardware forwarding.

Part of the reason you may be confused is that people are looking to build LAN's out of ATM switches as well. In this case, you could connect a local ATM switch to a Wide area ATM switch and have ATM directly between hosts attached to those switches. In this case, you wouldn't have routers in the loop at all.

A good way of looking at it is that if you are going from one type of network like ethernet to another net like FDDI, you need a router. If you are going from the same type of net to the same type of net, like ethernet to ethernet, you could use a bridge, or even a switch if you can build a switch to switch on the ethernet frame header.

Billing is totally orthogonal to the issue of routers and switches. We aren't being billed by volume, but simply by access class. So, 10 Mbps access costs X, 34 Mbps access costs Y, etc... You still have a DS-3 to your site from your supplier, but the switch will drop anything if you try and run at more than your access class for a given amount of time. This approach makes billing rather simple. You pay by the size of the pipe. It's possible to do this even if you have an DS-3 carrying both types of subscribers into the switch with ATM or something like SMDS. It's *very* hard to get an IP router to act like this.

the new cooperative agreement(s) are so limited, it is necessary to clarify these relationships and to indicate NSF's priorities and goals for NSFNET. This is especially important in the current atmosphere, which has increasingly and mistakenly assumed that the NSFNET effort is responsible for accomplishing broad data network connectivity for the nation's business and public sectors."

"NTTF believes that NSF should undertake a review, with community participation, of the appropriate number of access points to the high speed backbone, and accompanying increase in routing complexity, and establish criteria for evaluating proposals for establishing access points. The review should also address whether NSF can safely transfer the responsibility for funding and operation of the access points to the midlevel and private networks, thus eliminating the need for direct NSF involvement in this aspect of the evolution of the network."

The *COOK Report* notes that, if the NSF is to be responsive to its constituencies, it will have the conference committees and reviews they have requested. However, if it does have these meetings, it will have to work with uncommon speed to get the proposal awarded by the spring of 1993 and service begin in the second quarter of 1994. We expect that, caught between the figurative rock and hard place, as the NSF continues to thrash, the dates will continue to slip.

NTTF continues: "Several carrier representatives have suggested that the proposed operational and financial separation of backbone and NAP services will increase the total cost of NSFNET to the government and network users. In order to provide a basis for evaluating this issue, the solicitation should be revised to permit bidders on the cooperative agreement to propose both bundled and unbundled backbone and NAP services."

the NSF will be well advised to avoid any siren songs coming from the NTTF and others that would combine the provider of NAP and backbone services into one. Such a combination, when examined in view of the duties of the NAP provider to set tariffs and other connect charges for access to the commercial high speed backbone, would place a single provider of the high speed backbone and NAPs in a blatant conflict of interest position with the rest of the commercial community.

This may be true but the NSF will be well advised to avoid any siren songs coming from the NTTF and others that would combine the provider of NAP and backbone services into one. Such a combination, when examined in view of the duties of the NAP provider to set tariffs and other connect charges for access to the commercial high speed backbone, would place a single provider of the high speed backbone and NAPs in a blatant conflict of interest position with the rest of the commercial community.

A source sent the *COOK Report* a complete copy of the comments received by the NSF from 44 entities and totaling some 232 pages. Because it arrived shortly before press time I have not reviewed it extensively and will likely include more on it in the October issue.

A source sent the *COOK Report* a complete copy of the comments received by the NSF from 44 entities and totaling some 232 pages. Because it arrived shortly before press time I have not reviewed it extensively and will likely include more on it in the October issue.

A few highlights: Bellcore on behalf of itself and many of the RBOCs suggested that the NSF adopt a structure comprised of experimental and production networks . The experimental network(s) would support the Grand Challenges. The production networks would be an evolution of the current NSFnet three tier structure -- a praise worthy plan reminiscent of the editorial in the April issue of this newsletter.

Science International Applications Corp, (SIAC) sent a 22 page reply in which it outlined complete with cost estimates no less than three different NAP archictectures. It also included information stating that it had benchmarked the Cisco AGS+ Router at 60,000 packets or just over 30 megabytes per second. ANS sent 16 pages signed by Al Weis, Allan Baratz for IBM and Lawrence Bouman for MCI. MCI sent in 9 pages of its own comments separately signed by Diana Gowen. And NASA's 16 pages of comments signed by Milo Medin were quite impressive.

What is to Be Role of ATM and the IBM PARIS Switch?

According to a story in the August 17, 1992 *Network World*, many respondents to the NSFnet backbone draft solicitation emphasized their concern that ATM over SONET would not be a part of the NSFnet backbone solicitation. ATM (53 byte packets including 5 byte headers) is a CCITT international telecommunications standard that has been adopted almost universally by telecom vendors in Japan, the US and Europe. One would wonder why not? Is the NSF leaving the door open for someone?

A possible answer appeared in a July 27 *Network World* article by Ellen Messmer where she writes about trials of the Paris router with Rogers Cable Company in Toronto. According to Messmer: "IBM is preaping to test gigabit speed technology here that might become part of the new Internet backbone, if IBM catches the golden ring on the upcoming National Science Foundation contract. . . . IBM this September will test video conferencing and groupware applications over a a 6G bit per second backbone based on IBM's Paris routers. . . . . IBM [has also] developed ORBIT, a local area network technology which operates at 1.2G bit per second. The ORBIT and Paris protocol both use a proprietary protocol developed by IBM that can encapsulate the Transmission Control Protocol/Internet Protocol. . . . . To date one ORBIT ring has been attached to a Paris router [according to Alan Baratz who has overseen Paris design since 1988]. IBM hopes to have four Paris Rings operational by year end."

I wonder if PARIS qualifies as the oldest new technlogy on the Internet? IBM submitted the design philosophy to an obscure technical journal which in its publication of the resulting article in late 1988 market the article as received on May 4, 1988. Although IBM's Baratz apparently told *Network World* that PARIS is a router not a switch, PARIS has generally been spoken of as a switch and the *COOK*

Report has no reason to believe that it has suddenly been reengineered. It serves primarily to route level 1 and 2 using proprietary PTM over SONET instead of the internationally recognized CCITT standard ATM. In PARIS IBM has moved the proprietary nature of its networking from SNA at level 3 and 4 down to variable length cells (PTM) at level 2. And again, unless it has been re-engineered radically in the last six months, it encapsulates ATM cells within PTM - not TCP/IP within the proprietary protocol as Baratz apparently informed *Network World.*

In March of 1991 OTA was informed by a highly respected player in the design and construction of gigabit networks, that the fate of the PARIS switch would quite possibly determine whether the NREN was deployed as a public or private network. To put it mildly this heightened my interest in PARIS and what IBM's intentions appeared to be.

Besides, building private networks isn't really the goal of the NREN. We want to push technology in a direction that provides leadership for the national telecommunications infrastructure. We want to make sure what ends up getting built by the telcos meets the needs of data users, but the future is more than traditional packet services. It's going to be video, HDTV, voice and other multimedia environments, and it has to be supported by telcos, since they really are the only people who have the infrastructure to deploy these capabilities on a widespread basis.

NASA's Milo Medin declined to comment specifically on PARIS and PTM. However he did have the following to say on ATM: "The carriers have the bandwidth. If you buy the raw bandwidth and stick your own equipment on it rather than using the carriers' equipment, you will probably incur a cost hit because of this. Besides, building private networks isn't really the goal of the NREN. We want to push technology in a direction that provides leadership for the national telecommunications infrastructure. We want to make sure what ends up getting built by the telcos meets the needs of data users, but the future is more than traditional packet services. It's going to be video, HDTV, voice and other multimedia environments, and it has to be supported by telcos, since they really are the only people who have the infrastructure to deploy these capabilities on a widespread basis. That means you pretty much have to buy into at least their core architecture for broadband service provisioning, and that pretty much means ATM. ATM is far from perfect. But you have to look at it's cost/benefit tradeoffs and make a decision. We did that, and went with ATM."

On another but related subject Medin added, "we are very unhappy with any approach that locks in a given router or CSU/DSU vendor. We have significant leverage when multiple options are available. It keeps the vendors honest. Proprietary systems undermine that, and are really against the spirit of acquisition of open systems required by the government. Sometimes there just isn't any other way, so you have to hold your nose and go on. But it shouldn't be engineered into the system."

All this means that the United States Government should be very very careful that its major research data network doesn't inadvertently wind up embracing a switch that places the dominant high speed international data standard of the worldwide telephony community at a disadvantage. I argued this very forcefully within OTA during the spring and summer of 1991 where it was seen as a significant policy issue. Unfortunately I did not fully understand the implications of the NSF"s cost sharing philosophy at that time. I assumed that the bid would be judged primarily on technical excellence rather than cost and that the bidding process would show PTM for what it was, an unacceptable maverick.

However what has now become clear is that the NSF is intent on paying as little as possible for the new backbone and on giving the winning vendor the right to exploit fully the spoils in the market place. It has become very clear that IBM badly wants a high speed network on which PARIS (renamed PLANET almost a year ago) can be exploited. The *Network World* article is also important because Baratz choose to leak a statement that IBM may even submit its own bid for the backbone, outside of and apart from the one submitted by ANS. NSF bought a lemon with the T-3 network. With its cost sharing philosophy it may well buy another lemon with the backbone rebid.

The *Network World* article is also important because Baratz choose to leak a statement that IBM may even submit its own bid for the backbone, outside of and apart from the one submitted by ANS.

Jim Leighton, Manager of the Energy Sciences Network, (speaking for himself and not DOE) had the following comments in the August 17th *Network World*. According to Leighton, the NSF in its network plan "leaves many operational issues undiscussed and is potentially dangerous to the stability of the Internet." He said that NSF's demand for 155M bit/sec service is "unachievable in the near time and would not be deliverable to end users, even if it could be made available by a special implementation of facilities on the backbone." According to *Network World* Leighton "accused NSF of one upmanship with other Federal agencies for pushing for 155M bit per second. He said that no routers, interfaces, data service unit, channel service unit, or LANs supporting 155M bit per second currently exist. The NASA Science Internet office also questioned the need for 155M bit per second and, like Leighton, expressed the belief that the charge to connect to NAPS would result in a financial burden for federal agencies which now connect to the NSF backbone free of charge."

The Better Way of Backbone Selection

The August 17 selection by NASA and DOE of US Sprint as contractor for their respective increased high speed backbones showed a very different approach -- one based on technical excellence and close attention to stretching dollars to buy only needed services rather than hype. Milo Medin, an engineer and Deputy Project Manager of NASA Sciences Internet at NASA Ames, who was also involved in the the specification and procurement process has engaged in some detailed discussion with the *COOK Report*. This discussion forms

the basis for the concluding parts of this article.

Medin, while saying that he did not wish his remarks to be construed as criticism of the NSF, pointed out that "cost was not a factor in the selection That is, I and the others on the evaluation panel made our selection on the technical merit and project management only." He added that there was another area of management that did look at cost aspects with criteria to weed out of the selection process any unrealistic bids.

An interesting philosophy where the National Science Foundation with more money to spend takes the cost sharing "money saving" approach that emphasizes cost first and technical excellence second, while DOE and NASA with much less to spend emphasize technical excellence first and cost second and appear to have gotten the better deal.

While saying that he couldn't speak for the NSF's decision making processes, Medin said: "We have resources they don't. That is, the national labs and NASA field centers. We don't have to contract out the architecture and design of the network since we have internal expertise."

"It's important to realize there are real users who really need high speed networking to do their jobs. We have to provide it, but we have to make sure it really works. We want to do it in a way which agencies can share and cooperate properly, and act as a pathfinder in developing high speed infrastructure. Only 5% of the NASA HPCC program money goes to networking. If that networking doesn't work right, the other 95% of the money could be wasted. So we take this seriously."

"It's important to realize there are real users who really need high speed networking to do their jobs. We have to provide it, but we have to make sure it really works. We want to do it in a way which agencies can share and cooperate properly, and act as a pathfinder in developing high speed infrastructure. Only 5% of the NASA HPCC program money goes to networking. If that networking doesn't work right, the other 95% of the money could be wasted. So we take this seriously. We're on a really tight budget, so we have to work smart. There just isn't another way!"

"However, we also benefit from a strong focus on requirements, and we allocate budget accordingly. We have a fair amount of work to do. It's not easy work; but I think we gain a lot by leveraging carrier's work and work of the commercial router guys. Making something that runs fast isn't easy, but we really want to try and focus on a few new things, and use already working technology in areas where we can. Hopefully, we'll be successful. Having strong program and project management is a big help though."

Unlike the NSF way of doing business by cooperative agreement whereby it is generally unable to hold the awardee to any kind of strict performance standards, the NASA DOE deal with Sprint will be a five year contract for $50 million split between the two agencies in their purchase of services from Sprint.

Medin went on to point out that because NASA and DOE both have excellent technical network management and design skills, they are purchasing only level one and two services from Sprint. They will control routers and all other equipment at level 3 and 4 and be responsible for the TCP/IP service. "Performance requirements will be in the contract as well. We need these guarantees and certainly any commercial customer is going to require such guarantees before they buy such a service commercially. No one wants to be stuck with poor service and have no recourse."

While the August 24th *Network World* stated: "the DOE and NASA are expected [eventually] to order 155M and 622M bit service," Medin continued: "It may be profitable to elucidate the basic architecture of the NASA approach. IN FY 93, we expect to interconnect at high speed researchers at Ames Research Center, Goddard Space Flight Center, Jet Propulsion Labs, Langly Research Center, and Lewis Research Center using a high speed ATM switching fabric. But there are also requirements to support the awardees of 60 NASA Research Announcement Grants, at T1 rates. These will likely be distributed across existing NASA and NSFnet sites. The interconnects to the NSFNet are to support these aggregated T1 requirements, not T3 requirements. We will use the [Sprint provided] ATM network to interconnect the [aforementioned] centers to each other at DS-3 rates. The existing NSI and ESNET facilities already interconnect to NSFNet at the Federal Internet Exchange points at Ames and SURAnet in College Park. This architecture "will not change with ATM deployment, though obviously in the future, these interconnects will likely evolve."

Medin noted: " Also, both DOE and NASA have existing T1 network infrastructure that will remain in place until the DS-3 components are stable. Even after that point, we have many requirements that are only at T1 service levels. We will not be upgrading all of NSI to DS-3. That'd be silly. We don't have DS-3 or 622 requirements everywhere. We are a requirements driven organization. We only deploy service when documented requirements dictate it. Because our focus is much narrower than NSF's, we have to support a narrower set of requirements as outlined in the HPCC program. Since the ATM network will be plugged into NASA routers that run Open Shortest Path First (OSPF), we expect to be able to switch back to the T1 path in less than 10 seconds. We routinely see that level of route switching in the T1 network today. We expect to be able to handle that in the ATM network as well. It is not organized as a separate T1 and T3 network as is the case with the NSFnet backbone. The government operates both parts so there is no need to split them that way in our case. There will also be a phase in period where the ATM network operates as secondary to the T1's. We won't throw the switch until we're sure it will work."

We can get very fine grained service increments. If we bring up a site at 10 Mbps service class, and it's not adequate, we can order up 15 Mpbs, and Sprint can send a management command to the ATM network and implement it. We have a diverse set of applications that the network has to support. . . . . We look at growth and make sure the network can evolve to meet existing and new requirements. Even if we could tech-

nically deliver gigabits to everyone, it would be unlikely that we could afford it. So you have to prioritize.

Asked if this meant that NASA acknowledged that everyone did not need gigabits and was building a flexible network architecture to deliver high bandwidth on demand only when really needed, Milo stated: "The ATM service will be able to provide a variety of service classes, up to the speed of the local loop to the ATM network. We can get very fine grained service increments. If we bring up a site at 10 Mbps service class, and it's not adequate, we can order up 15 Mpbs, and Sprint can send a management command to the ATM network and implement it. We have a diverse set of applications that the network has to support. They have different requirements for capacity and latency and such. We look at the requirement, and future growth potential, and implement to specification. That's just engineering. We look at growth and make sure the network can evolve to meet existing and new requirements. Even if we could technically deliver gigabits to everyone, it would be unlikely that we could afford it. So you have to prioritize. And that means you need to understand the requirements that are being supported."

"There are also other aspects to the procurement. Support for SMDS interoperibility, integrated broadband video support, priority and resource control mechanisms, and a local SONET testbed in the bay area to prove out technology and try things out before they go into the production network."

Medin concluded: "Also, note that NASA and DOE operate COTS type networks (COTS = commercial off the shelf). No components of the Sprint system are being built especially for us. All will be available to others to buy separately or in a system. Only doing things in this way minimizes government investment in product development. We don't have big budgets in NASA or for high speed networks. We gain the most leverage for our program dollars by working in directions that the carriers already are going in, but serving as an incentive for them to accelerate progress in these areas. We want them to be able to take this technology and transfer it into the commercial sector where others can procure it as required. Only in that way can you truly build a national infrastructure. The Government doesn't have the resources to do this by itself."

We want them [the carriers] to be able to take this technology and transfer it into the commercial sector where others can procure it as required. Only in that way can you truly build a national infrastructure.

Conclusion

It seems to the *COOK Report* that the backbone design choice may be considered a matter of philosophy. The NSF, with its larger budgets could hire a network design staff, if it so choose. The NSF does not, repeat . . . does not need to use its NREN funds to push gigabit technologies. The carriers with ATM and SONET are doing quite an adequate job of pushing gigabits on their own. The NSF, in ordering under a "toothless" cooperative agreement, a 155M bit per second backbone to connect everyone and everything now with no true indication of need is - if past IBM performance is likely to be any guide - is unlikely to push any envelopes any further or faster than either NSI or ESnets. Furthermore, it may give inadvertent help to the spread of proprietary protocols if PARIS & IBM "catch the golden ring." It is well know within the community that Alan Weis and Alan Baratz have been close colleagues. With Baratz volunteering to the trade press that IBM is considering bidding without ANS, ANS should indeed be concerned about its future. In the mean time staff people on the Hill should stand less enamored of their gigabits at any price and pay attention to the issues raised in this article and in discussions else where in the Inernet and trade press.

DOE Announces Innovative "Grand Challenge" Supercomputing Projects

On August 14, 1992 the Department of Energy (DOE) announced six new projects under the auspices of the HPCC initiative. Massively parallel supercomputers will be used by researchers to model groundwater pollution, enhance oil recovery, burn fuel more efficiently, and better understand fundamental issues in human health through collaborative research grants.

The six projects will receive a total of $20 million over the next five years, including an estimated $8 million in cost-sharing by universities, states, industrial companies and foundations. A total of nine laboratories, 16 universities and more than 40 businesses will participate in or fund the projects.

The DOE is the first agency to announce its awards for Grand Challenge projects. The six projects were selected from a field of 48 proposals through a peer review process, in addition to a review by an inter- and intra-agency Grand Challenge Review Panel.

While the research teams involved in the six projects are spread across the nation, they will use the National Research and Education Network to communicate with each other and with the High-Performance Computing Research Centers recently established by the DOE. These centers, located at Los Alamos National Laboratory and at Oak Ridge National Laboratory, will provide access to the massively parallel machines.

The six Grand Challenge projects involve: computational chemistry, structural biology, fluid and combustion dynamics, petroleum reservoir and groundwater modeling, fusion energy, and high-energy physics. Some specifics:

Computational Chemistry: This project will study the incineration of chlorinated hydrocarbons. Developing better ways to incinerate these chemicals will assist in clearing up toxic waste sites.The project will also study the development of replacements for chloro-fluorocarbons, chemicals implicated in the destruction of the ozone layer in the upper atmosphere. Additional areas of focus will include reactions of pollutants on and in clay minerals, processes important in the reclamtion of contaminated soils and groundwater; and ways to redesign natural enzymes to better degrade pollutants in soil and groundwater.

The project's major participants are Argonne National Laboratory, Pacific Northwest Laboratory, Allied Chemical, E.I. duPont, Engelhard, Exxon Research and Engineering, and Phillips Research Center.

A Structural Biology Project will develop and use software to study proteins, protein "folding" (structure) and protein behavior, and detailed DNA sequence analysis and interpretation. Proteins are key to nearly all body functions: growth, digestion, healing, fighting diseases, storing and releasing energy, and degradation. Better understanding of DNA will be important to grasping the implications of hereditary characteristics, including diseases and susceptibility to disease.

The California Institute of Technology and Argonne National Laboratory are the major partners in developing the enabling software components and computational resources for a collaboration that involves the University of California at Los Angeles, Yale University, the University of Washington-Seattle, and J.D. Searle.

Computational Fluid And Combustion Dynamics: The major purpose of this project is the development of new methods for computational fluid dynamics and apply them to the problems of simulating combustion dynamics. Computer simulations will be used to develop more efficient and less polluting power-plant burners, waste incinerators, and internal combustion engines. . The enabling mathematical and computational components will be produced by the University of California at Berkeley, Lawrence Livermore National Laboratory, New York University (Courant Institute of Mathematical Sciences), Los Alamos National Laboratory, and Sandia National Laboratories at Livermore.

Petroleum Reservoir And Groundwater Modeling: The object of this Grand Challenge project is to develop the algorithms for modeling flow through porous geologic formations. This will enable more effective studies of groundwater pollution migration and reclamation and enhanced oil recovery from reservoirs. Of course all this will be done with the goal of making the processes more economical. The project is a collaboration between the Center for Petroleum and Geosystems Engineering at the University of Texas at Austin (oil recovery) and the Center for Research on Parallel Computation at Rice University (groundwater transport).

Numerical Tokamak: Tokamaks are doughnut-shaped fusion reactors, and are at the threshold of demonstrating the long-sought goal of scientific 'break even' -- producing as much energy as it takes to run the experiment.

Through this project, high performance computing will play a major role in designing future machines. By using better computer simulations of complex plasma physics and mechanical design, the overall development costs of a fusion reactor should be reduced. The research consortium for this project involves scientists and engineers from Los Alamos National Laboratory, Lawrence Livermore National Laboratory, Oak Ridge National Laboratory, Princeton Plasma Physics Laboratory, General Atomics Inc., Jet Propulsion Laboratory, University of California at Los Angeles, and the University of Texas at Austin.

High-energy Physics: The search for an understanding of matter at the subnuclear level has long been one of the biggest users of computer power. This project will use the new massively parallel computer systems to improve our understanding of the basic properties of matter by making predictions that will allow more precise experimental tests of the Standard Model of fundamental particle interactions.

The project team has members from the University of California at Santa Barbara, Los Alamos National Laboratory, the University of Arizona, the University of Washington-Seattle, Oak Ridge National Laboratory, the University of Washington-St. Louis, the University of California at San Diego, Thinking Machines Corp., Ohio State University, Indiana University, Syracuse University, the University of Colorado, and the University of Utah.

NLM's Director Appointed HPCC Coordinator: Substance or Window Dressing?

According to the August 24, 1992 *Federal Computing Week*: "The White House Office of Science and Technology Policy has upped the ante on its High Performance Computing and Communications Program by naming a key administration official to lead the five year, $1.9 billion interagency effort. Dr. Donald Lindberg, Director of the National Library of Medicine, was named director of the National Coordination Office for High Performance Computing and Communication. In the post, Lindberg will serve as a liason to participating agencies, reporting to White House Science advisor Dr Allan Bromley."

Several phone calls by the *COOK Report* to sources at the National Library of Medicine revealed that Dr. Lindberg will remain in his post as Director of the NLM. It is unofficially estimated that he will spend 40% of his time as Director of the HPCC Coordinating Office and 60% as Director of NLM. His full salary will continue to be paid by the National Institutes of Health, so it should not be considered a political appointment. The appointment does in fact run for two years. The August 19th NIH press release states that "the establishment of a national coordination office recognizes the need not only as single point of coordination as HPCC progresses in its implementation, but also a primary interface to Congress, academia and the industrial sector." *Federal Computing Week* commented that the program's management had been "handled in an ad hoc manner by the FCCSET" chaired by the NSF's Director Walter Massey. Interestingly enough FCCSET also reports directly to Bromley.

A slap at the National Science Foundation for the controversy surrounding the NSFnet back bone? Perhaps. What is uncertain is what real powers that Lindberg may have that others don't. We wonder how anyone spending less than 100% of his time could stand a chance of co-ordinating this NREN HPCC beast that still have no real operational plan. We found it interesting that the information officer at NLM couldn't tell us where in the Federal Science bureaucracy the new office would be located and that Dr. Lindberg would soon be leaving for two weeks in Switzerland. We hope that this is more than window dressing but will believe it when we see it.

At Play in the Field of the Boards: Report on ONE BBSCON

by Steve Cisler

From the Editor - Thanks to Steve Cisler for permission to reprint.

Denver, Colorado: August 13-16, 1992.

Our plane punched down through the low clouds a half hour late and landed at Stapleton airport, just minutes from the Stouffer Hotel where all the activities for the Online Networking Exposition and BBS Convention were to take place. This was to be the first ecumenical gathering of bulletin board system operators (sysops), hardware and software vendors, and programmers that spanned the DOS, Unix, and Macintosh worlds. There had been annual meetings of Fidonet sysops, but the parent organization had folded even as the number of BBS systems continued to grow.

Jack Rickard, President of ONE, Inc. and publisher of *Boardwatch Magazine* <jack.rickard@boardwatch.com> partnered up with Phil Becker, author of TBBS software, to organize a trade show for BBS operators that would be inexpensive enough to attract those running their boards as hobbies but with tutorials and panels with subjects that would attract the entrepreneurs who are working from their homes or small offices and those who have mounted and maintained successful business systems with four, eight, on up to 64 phone lines for their clients.

Most of us in the online world are fairly stratified in our interests when it comes to networking and communications. There are MIS shops depending on minis and mainframes; there are the academic and research networks where NREN and the Internet is of paramount interest; there are consumer services which have been a continued failure when promoted and maintained by the regional phone companies and have been a mixed bag when we look at Prodigy, GEnie, CompuServe, and America Online.

The online industry of BRS, Dialog, and Mead Data is trying to break out of its own mold and attract more than librarians and devoted professionals who need the high priced information and are willing to learn to tolerate difficult interfaces.

The world of BBS systems, users, and sysops has never been validated by some of the mainstream opinion makers. Many of the operators and users have been outsiders, socially, politically, and even economically.

ONE BBSCON brought many of these outsiders together for the first time, but it also attracted mainstream users who have found BBSes to be cost effective, easy to maintain (compared to other sorts of electronic systems), and extremely useful. The show put the spotlight on the "Industry" as Rickard and Becker hope it will become. BBS System vendors were on the organizing board, and there were tracks for the major DOS/Intel systems: Wildcat!, TBBS, PCBoard, MAJOR, and Searchlight as well as other operating systems such as Unix and Macintosh (the latter comprised less than 5% of the attendees). Other tracks were for legal and social issues, corporate and business applications, how to make money with a BBS, the Internet/NREN, Mail networks, and a technical track.

The opening reception included just a sprinkling of ties, making me feel over-dressed in a sport coat. With the dress code set for the rest of the day ahead, I thumbed through the program and spoke with BBS and networking folks whom I had met online or at other conferences. Dave Hughes of Old Colorado City Communications <dave@well.sf.ca.us> and Frank Odasz of Big Sky Telegraph <frank@bigsky.dillon.mt.us>. I noted the lack of education tracks in the program, but with the participation of these two pioneers in the program education became an important subtheme. The range of activities people are involved in is indicative of the power and freedom to experiment that goes with cheap hardware, BBS software, and a reasonably priced public telephone network. A journalist from Albuquerque is running his paper's 16 line system which provide news to callers at no charge. Another team had just returned from Russia where they were helping set up a country-wide system using low-orbiting satellites for data transmission. An entrepreneur who had been sued by Playboy enterprises for vending GIF images of Playmates (and using that trademarked name in his BBS menus) without much thought of copyright was present.

Day One

At the opening session it was evident why this was going to be a successful conference and why BBSes would become an industry that would overshadow but not eliminate hobby uses of such systems: cheap 9600 bps modems, inexpensive 486 servers, and telephone systems that are not charging by measured use (as is done in many other countries).

Phil Becker made a strong pitch for the BBS as a business tool for mainstream activities and belittled what he called the 'stupid niches' like Keith Wade's *The Anarchist Guide to BBS*. Clearly, Becker wanted to encourage mainstream business uses of this technology. They were impressive and diverse. Others like Tom Jennings (author of the extremely cheap and popular Fido software) sees the technology benefiting the outsiders: the fat, the handicapped, the socially inept, the disenfranchised, the radical, the non-mainstream. And of course, BBS technology can fill both Becker's and Jennings' very different visions, but Becker's was the one emphasized here.

What I want to continue to emphasize in this report is the diversity of users, of models, of software, and of business models. *Boardwatch* estimates that there are 44,000 public systems (and many more private and corporate ones) in the U.S. The four vendors sit-

ting on ONE, Inc. board claim to have an installed base of 50,000, and this excludes Fido, the most popular of all systems.

The Keynote speech was given by John Dvorak who writes for many computer magazines, has co-authored a successful book on telecom, and is often on the public speaker circuit. His columns are entertaining, opinionated and as Art Kleiner said in an old S.F. Bay Guardian article, he is a 'curmudgeon without a cause.' People who like to rattle cages can be good speakers and warm the crowd up for the ensuing events. However, Dvorak devoted so much time to self advertisements and plugging his books that it seemed he must have spent about 10 minutes thinking about some of the issues that needed addressing with regards to BBSes.

A sprinkling of Dvorak comments: "get a fan for your 486 machines... OS/2 is fun! (Windows is not)...the BBS community needs a lobbyist in each state; it's embarrassingly naive and should examine how it operates on different levels. Al Gore is the Dan Quayle of the Democratic Party. The porno boards are always under scrutiny by the govt. one way or another... Playboy images on a BBS has to be called fair use... BBSes cannot continue to allow slander on the boards. You have to clean up your act by self-policing.

He also called for a constitutional amendment to protect electronic rights, but he did not know that constitutional scholar Laurence Tribe had proposed this in his keynote at Computers, Freedom, and Privacy 1 eighteen months ago. On the one hand Dvorak pandered to anti-Congress sentiments in the crowd and then said to get involved in politics; yet he found the Perot candidacy a sham, including the electronic element.

A Few of the One Hundred + Panels and Presentations

I attended the Electronic Frontier Foundation program which filled up with an overflow from the Legal track where Lance Rose addressed many questions about the rights and responsibilities of sysops in an increasingly litigious and regulated environment. Shari Steel, a lawyer in the EFF Washington office, explained their activities and then answered questions which revealed the strong anti-Washington, anti-lawyer, anti regulation sentiments in the audience. It is evident that many sysops value the control they believe they have when they set up and run a bulletin board. Hearing about the FCC and Congressional efforts to change 'their' world made some people angry and made others want to organize. Steel handled their questions honestly and admitted when she did not have clear cut answers.

Midway through the conference a group of software developers decided to organize to try and learn about existing standards before setting ones that would benefit their own developers and users. The three main areas include interface, messaging and document structure, and graphics. Surprisingly, some did not want to give up the diversity that is so evident in the many interfaces and message protocols, but most agreed to try and set an agenda via electronic mail and plan for other actions at the next conference. The Internet Engineering Task Force use of RFCs (Request for Comment) was held up as a model they would emulate. Few had patience and resources to set up a very slow-moving mechanism such as the ISO committees.

Jim Warren moderated the panel on electronic democracy where most members were excited about the power of the tools. Gary Stryker of Galacticom proposed a system called Super-Democracy which would include continuous electronic voting on issues in a hierarchy by geographical region. Shari Steel reminded the group how many did not have computers or modems, and that electronic democracy would exclude many potential voters. Gary Nakarado, a PUC commissioner in Colorado started a BBS to learn more about the medium and to be in touch with the interesting people in his community. Unfortunately, they have not been logging in. PUC activities attract very little attention and he has very few calls and questions from the general public. He is interested in having more input on issues such as ISDN service, BBSes, and other issues affecting Colorado utility users.

Bernard Aboba, author of *BBSes and Beyond*, talked about the software for connecting Macs to the global mail networks (RIME, Fido, Internet, uucp) and it was evident that the Mac is a terrific front end for many systems, but as a server it needs more power and more tools from third parties.

Developers from ResNova (714) 840 6082 showed how their Mac fBBS software could fill the gaps, as did a rep. from SoftArc (416) 299 4723. SoftArc's FirstClass server and client software looked very powerful and full of features that would allow a FirstClass BBS to serve many concurrent users on LAN, dialup and TCP/IP access. All of the companies are quite small, and the wish lists of new features grows faster than the staff to work on them, but I was amazed at the power and sophistication of the DOS and Mac BBS systems. Event Horizons, the BBS vending adult GIF images south of Portland, Oregon, has a 64 line 80486 system running on TBBS! Other systems running multiple cpus have a hundred or more lines coming in. Clearly, these are not basement run, part time operations.

Diversity

One morning I went down to comb the literature tables and read the cork bulletin board where a variety of fascinating notes had been posted. They will give you a sense of the diversity in this community:

-Monterey Gaming System 408 655 5555 (free)

-Black Cat Information Service in Rochester NY 716 262 3680 (Visa/MC accepted!) Games, Society for Creative Anachronism files and Adult Info

-the Zoo...an electric safari. your tour guide: Chuck 2. 312 907 1831 to 183

-The OU BBS, University of Oklahoma(telnet oubbs,telecom.uoknor edu)

-Power Windows! BBS (also for OS/2 users) Huntsville, AL 205 881 8619

-The Invention Factory, (NY, NY) 212 274 8110

-The Online Diver (Brooklyn Center, MN) 566 5267 No area code must indicate that it's for local Minnesota divers primarily.

-Nautilus Commercial Data System with 250,000 public domain files, 200

incoming data lines (!), 28 gigabytes of storage, satellite weather images, hourly news updates, games, dBase templates, GIF images---all out of Iola, Kansas 316 365 7631.

-Infinity Complex "a wickedly addictive Science Fiction game for MAJOR BBS systems. Infinity Complex puts your users in a bizarre arena of the future, where they must battle for their very lives...and use up a lot of online hours in the process!" 403 476 8369 (voice)

-an ad for the first annual Puget Sound (WA) BBS convention (no phone contact)

-ads for serial port boards, new BBS software, consulting services, and calls for source code for data compression.

-BAWIT Bay Area Women in Telecom for working women in telecommunications in the San Francisco area. Contact bawitrequest@igc.apc.org.

-Make your own custom CD-ROM for $199. Up to 640 mb. ISO 9660. 800 762.

Internet and NREN

There was a lot of interest in Internet/ NREN issues, but only a few people knew much about them. The panels on Internet connectivity, legislation, and interfaces drew good sized crowds but needed more basic information in a standard presentation format before having Q&A. BBSes can be a good interface for people going on to the Internet. It provides a way of formatting and filtering the anarchy of the Internet, even as it offers occasion for excess control of what a caller can see and use. I spent more time showing resources using Mac-based interfaces than talking about the intricacies of the growth of NREN in my session which was included in the small Macintosh track rather than the larger Internet track. I also participated in a graphics discussion where the panelists discussed GIF (the CompuServe standard so popular on BBSes and the Internet), NAPLPS (which is good for multi-lingual communications and small vector-based images), JPEG (the compression du jour that may displace GIF and the one that the Smithsonian and Apple are touting for Project Chapman), and FIF (fractal image format which is a more efficient proprietary algorithm than JPEG but which takes a long time to compress).

Summary

The BBS world is changing, growing, exploding. Jack Rickard has provided good coverage in his magazine, and his conference was a big success and and a very good value considering the amount of fine food that was included with the conference activities. If you are not in the BBS world, and even if you are, it's hard to be aware of the activity because it is so distributed. This conference helped immensely.

I think that Rickard will have to face a problem of success: will he continue to be the lively and opinionated journalist when his magazine and his conference become the focus of the whole industry and a possible industry association? He may have to defend actions when he should be exposing them, but that is looking a couple of years ahead. Right now, there is no way to go but up and out because of the growing interest in this medium of information dissemination and of personal communication. A BBS provides both sysops and users an enormous amount of leverage, and the library world should take notice more than it does. One prominent public librarian who is quite involved in electronic dissemination of information remarked to me a few years ago that it would be great if BBSes just went away. I have heard other dismissive or even snobbish comments about BBSes, but the four librarians whom I met at ONE BBSCON all realized this is foolish. It's not the only tool to use, but it can be a very important one.

Contacts mentioned in the text: people and products

Gary Nakarado, Colorado Public Utilities commissioner <nakarado @well.sf.ca.us> (303) 526 5505 is his BBS number NAPLPS: North American Presentation Level Protocol Syntax. Dave Hughes <dave@oldcolo. com or dave@well.sf.ca.us> Fractal compression: Fracterm, Inc. in Richmond, BC (800) 676-3111 GIF: CompuServe Art Gallery and various browsers at ftp sites JPEG: mail jpeg-info@uunet.uu.net or contact sac@apple.com *Boardwatch Magazine*: (303) 973 6038 or jack.rickard@boardwatch.com TBBS software: eSoft (303) 699 6565 Super-Democracy Foundation: (305) 583 5990 Bernard Aboba: BMUG, Inc., (510) 547 0345 Jim Warren (electronic democracy): jwarren@well.sf.ca.us Electronic Frontier Foundation: eff@eff.org Lance Rose (legal issues): author of Syslaw a legal guide to the rights and responsibilities of sysops. Laurence H. Tribe, "The Constitution in Cyberspace" anonymous ftp from ftp.apple.com in the /ftp/alug/ rights directory 49 kb.

Public Libraries and the INTERNET/NREN: New Challenges, New Opportunities

On July 1, 1992 Chuck McClure, Joe Ryan, Diana Lauterbach, and William Moen of the School of Information Studies Syracuse University published a useful report titled *Public Libraries and the INTERNET/NREN: New Challenges, New Opportunities.* As part of a year long study involving interviews and data collection from public libraries.

Some of the study findings:

While public librarians are generally enthusiastic about the prospects of national networking, they are still not quite sure what it has to offer their institutions. The librarians suffered from the conceptual gap of hearing about networking without generally ever having *experienced* it.

Public Librarians felt that inadequate attention to NREN had been given in their literature. Most were surprised to find out that it could be useful to them.

They expressed concern about taking technology risks, because of the poor economic implications of the current climate. Nevertheless they felt that public libraries should not become completely averse to risk taking.

A barrier to internet connection for public libraries is the dearth of information on how one does get connect

(continued bottom of next page)

Report from ONE BBSCON

by Dave Hughes

ONE BBSCON was probably the most significant gathering of 'bbs and beyond' practitioners yet held in the world. Over 1,000 attendees, probably 900 of whom paid to get in. Four solid days of 109 separate sessions, an active 50+ vendor area where specific (I am not at liberty to name them) vendors of BBS and related hardware and software grossed up to $30,000 in sales to the attendees - at least 700 of whom operate some BBS-type service. (This was determined when a speaker got all those who ran systems to stand up during the plenary session, and sit down based on year they started their first system, until only Ward Christensen was left (1977). And lots of computer, communications, as well as some general national press.

The event was organized to move 'beyond' hobby BBS conferences, to the serious commercial, professional, level. A break with the last 'Fidocon' conference of the year before, which drew 400+, among whom were business, government, education, professional people who wanted more than a chance to hob nob with the 11 Fido software developers. They flocked to BBSCON, held at a major Denver airport hotel which sold out as did overflow hotels. Substantial participation by foreign nation attendees. In one session I chaired, 15 nations - from Slovinia to Mexico - were in the audience, as well as the usual suspects. At least 8 came from Sweden - and from the Datapac national level, not just the BBS level.

I was so busy either participating, chairing, or presenting sessions (on international, educational, Naplps, graphics, ISDN, telco relations) I did not get to attend many sessions, where hundreds jammed in, and could barely get around the vendor area one full time. (I also had commitments to present across town to a far less attended successor to the Fidocon-91 conference, so shuttled between - about 5 times. And where three radio programs emanated from. The most noteworthy event of which was my first squaring off with Jerry Pournalle in the plenary session over electronic politics among other things, but followed by a dinner in which we swapped Korean War stories, and ended with Jerry showing unusual interest in my work with Naplps, the Indians, Russian software engineers, and global any-font telecom. During an hour and a half in his suite, swapping disks on notebook computers)

Whether any 'notables' were interested in Naplps or not was of little concern to me, because it was abundantly evident that Naplps, a year after I was able to release the first shareware drawing program, has started to take off. With Phil Becker, eSoftn TBBS vendor, telling me hundreds of TBBS owners are implementing a $199 program which incorporates it. Ditto from Jack Rickard. And numerous queries and statements from other BBS developers who haven either partially, or fully implemented the standard in their 'boards.' Including some Europeans who like the ability of Naplps to support their language fonts. With the President of Microstar flying in from Ottowa, appearing on panels, and reporting many queries from Europe, and making deals right and left. And even Norpak showing up wondering what all the fuss was. And with R&D heads attending from Hayes (Smartcom) and Datastorm (Procomm) approaching me to find out what they have to do to support Naplps.

I just hope progress continues with our software engineers in Russia toward Old Colorado City's own program suite in Naplps, so that we can make some sales too. I got 7 years and lots of money to recover from trying to sell my conviction that Naplps is the most suitable low bandwidth telecommunications graphics standard.

There was, lots of interest in 'graphics' in general, (Publisher Rickard says it is a fast rising demand) and while BBS operators incorporating GIF, especially for chat lines and making money hand over fist at it were popular, there were plenty of other graphics approaches present and attracting attention, from Coconet, Apple Mac, Resnova, Jpeg and others.

All in all BBSCON was a smashing success. And next year will be two to three times as big, and twice as profitable for everyone.

---

ed. Generally the librarians wanted to know about costs and how to get connections and to know *now!*

During some of the focus groups held by the researchers, some of the public librarians began to see that the internet could be a two edged sword for the public library. On the one hand they thought that public access terminals in the library were good ideas. But, on they other hand, they got exra food for thought when they realized that those with computers and modems could access the information on their own and might never need to come to the public library for *that* purpose. Librarians expressed interest in some of the following services:

full text and color children's books on the network

practical news groups or listservs such as auto-repair, recipes, homework tips, crafts, etc.

community based information services in health care, community activities, and unique local resources.

a job network

and remote access to library reference services and materials.

Copies of this very well done report may be obtained for $15.00 (includes postage & handling) from the Publications Office, School of Information Studies, Syracuse University, Syracuse, NY 13244-4100.

Router Flapping Problems on the T-3 Backbone

Editor's Note: An acquaintance forwarded to me the following T-3 trouble report from MERIT.

From: "Mark Knopper" <mak@merit.edu>

Date: Tue, 11 Aug 92 22:47:35 EDT

Subject: EGP-related routing problem on T3 backbone

During the last 72 hours, we have seen a multiple instances of routing instability across the T3 network including RS6000 CPU starvation on several ENSS nodes and route flapping on E139, E138, E136, and E134. Several sites have called to complain about route flapping and instability. We spent considerable time over the last couple of days trying to isolate this. The following is our current understanding of the problem, and our plan for correction. We will update the list as we learn more.

With the new network announcements that resulted from the scheduled configuration update last Friday (8/7), the T3 ENSS nodes started sending EGP updates to regional peer routers that exceed 8KB in size at all sites configured for explicit routing. The problem starts when the ENSS sends the regional peer an 8KB+ update. The peer router may flap if it is operating with software that will NOT support 8KB+ EGP updates.

Also, there is a dormant bug in the RS6000 rcp_routed EGP code which involves routes getting imported via IBGP which do not get flushed out of a queue. When the regional router flaps (misses one message from the peer) due the 8KB+ update described above, the rcp_routed routes derived from EGP sitting in the queue do not get installed and this might indirectly result in a routing inconsistency between the ENSS and its CNSS neighbor.

We suspect that a simple fix to the T3 network problem is to flush the queue every time we timeout the EGP derived routes. We have a new version of rcp_routed that flushes the queue and has a trace statement that logs the event. We will install the new rcp_routed on C99, C51 and if successful, we will install it on E138 at 5am EST 8/11 (with SURAnet approval). If this is successful, we will install it on several other nodes as emergency maintainence on 8/12pm. We will send another note to the nwg list tomorrow with an update to the problem.

However this fix not solve the problem of the 8KB+ updates causing route flapping on several regional peer routers. Peer networks that are running BGP should not experience this problem. We have contacted Cisco, Proteon, and Wellfleet, and have learned the following regarding their suggested software fixes to this problem.

Cisco

- -----

Experiencing this problem depends on the version number of software that you are running. To find out what you're currently running at, do a

"show buffers" and note the size of the huge buffers.This is the maximum size IP packet that the router can reassemble. If EGP updates come in that are larger than this, then you will get reassembly failures which can be seen in "show ip traffic".

In later Cisco releases, there is now a knob so that you can change the buffer size on huge buffers. Using this, you can reassemble up to 64k IP packets. The following releases support the following buffer sizes:

v8.1 8KB buffer
v8.2 12KB buffer
v8.3(1) 12KB buffer
v8.3(>=3) 18KB buffer, but configurable
v9.0 18KB buffer, but configurable

Proteon Routers

The Proteon router has a fixed size reassembly buffer. Any packets bigger than the reassembly buffer will be dropped. Proteon will generally increase the size of the reassembly buffer in each release. Its current size (in Proteon releases 11.0 and greater) is 12K. In release 10.0b, a large number of which are probably still in the field, the reassembly size was 8K.

If there is a site that thinks it is having this problem with a Proteon router, they can contact Proteon customer service to get the latest software revision. Customer service is familiar with the reassembly buffer issues.

Wellfleet

- ---------

The Wellfleet router also has a fixed size reassembly buffer. Any packets bigger than the reassembly buffer will be dropped. Any Wellfleet router running a software release of v5.6 or older will have a 4KB buffer. This release stopped shipping 2 years ago. All new releases are compiled with a buffer size of 16KB and should not experience this problem.

- --Jordan Becker, ANS
Mark Knopper, Merit

IBM and SEARS Form Giant Global Network Outsourcing Company

Shades of Prodigy!? Or is networking becoming ever more critical to IBM? The August 24th issue of *Network World* announced as its lead article the formation by IBM and Sears of Advantis, a huge global network outsourcing company that consolidates the VANs (value added networks) of the two founders.

According to *Network World*: "From day one Advantis will be a giant company. It will handle the internal networking needs of IBM and Sears while serving more than 9,000 customers worldwide. Analysts expect Advantis to generate $1.1 billion of revenue during its first year of operation."

"The Advantis network will combine IBM and Sears high speed Systems Network Architecture backbones and support leased line access at speeds up to T-1 and dial access at speeds of 300k bit to 19.2k bit/sec. Access to the net will be from 550 cities in 92 countries worldwide, either directly or through Advantis affilaites.

Adavantis will offeramong other services: SNA, X.25, TCP/IP, APPN, ... and a range of messaging and EDI services."

INTERNET Domain Survey Reveals Continued Explosive Growth In Network

Network Information Systems Center July 1992
SRI International
Internet DomainSurvey

The Domain Survey attempts to discover every host on the Internet by doing a complete search of the Domain Name System. The latest results gathered during late July 1992 are listed.

-- Mark K. Lottor

Hosts: 992,000 [890,000 hosts in April 1992]
Domains: 16,300 [approximately]

Network number analysis of all IP addresses listed in the DNS:

Percent Of Network-Address-Space Capacity Being Used

Number of Networks	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
Class A: 60	100%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Class B: 2714	99.9%	0.1%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
Class C: 3795	85.0%	7.6%	3.0%	1.5%	0.9%	0.9%	0.4%	0.2%	0.2%	0.0%
Total: 6569										

[for example: 3.0% of class C networks have 30% of the available class C address space (256 hosts) assigned to hosts]

Host Distribution by Top-Level Domain Name

326630 edu	21021 se	2902 kr	943 br	74 hu
277551 com	19192 fr	2803 il	789 mx	50 int
62584 gov	17188 ch	2733 dk	633 pl	43 ee
48639 au	15757 jp	2410 hk	624 ie	14 yu
43907 de	15718 fi	2118 tw	616 gr	8 tn
38929 ca	14354 no	2073 za	569 cs	6 in
37776 uk	7044 net	1831 nz	400 is	5 ve
33161 mil	6489 at	1532 be	201 us	5 aq
25896 org	5147 it	1318 pt	165 cl	2 ar
21105 nl	3603 es	949 sg	80 lu	1 th

Top 50 Host Names

508 venus	377 gw	303 eagle	263 fred	240 mac12
471 pluto	360 mercury	301 mac5	262 sirius	240 mac11
438 mars	355 iris	296 gauss	262 mac8	236 alpha
421 cisco	351 pc2	292 pc3	262 hermes	235 mac13
408 pc1	337 mac3	284 mac6	259 mac9	234 mac14
407 jupiter	334 orion	275 hobbes	257 calvin	231 phoenix
405 zeus	320 mac4	274 mac10	254 pc4	231mozart
405 mac1	310 newton	271 mac7	253 apollo	231mac15
383 mac2	310 neptune	266 thor	248 athena	227earth
381 saturn	306 charon	263 merlin	243 titan	226mac16

Editor's Comments: one should note the amazing growth of commercial or .com domains. They will soon over take education (.edu) domains in numbers. Mil means military and is used by milnet. Org means organization as for example a foundation or non profit or consortium. Net means network. Most of the other domains stand for country names. Au, I believe, is australia and de is Germany. UK is United Kingdom and after org in descending order: Netherlands, Sweden, France, Switzerland, Japan, Finland, Norway and so on.

From CNRI & IETF: An Internet Service Provider List

Editor's Note: Here from the Internet Engineering Task Force Mailing List is a recent list of internet service providers. It is not comprehensive, but is more complete than anything else we have seen.

General Assistance

For assistance in determining where and how to connect to the Internet, send email to:

1. **MERIT** (University of Michigan)
   nsfnet-info@merit.edu

   or call: +1 800-66-MERIT (+1 800-666-3748)

2. **FARNET**
   Ms. Laura Breeden
   Executive Director
   FARNET, Inc.
   Work Address: 100 5th Avenue
   4th Floor
   Waltham, MA 02154
   USA
   Work Phone: +1 617-890-5117
   Work Fax: +1 617-890-5117
   Work Phone: +1 800-723-2763
   Email: breeden@farnet.org

3. **NSFNet Network Service Center** (NNSC)

Ms. Corinne Carroll, User Services
Work Address:
NSF Network Service Center
(NNSC) Bolt Beranek and Newman, Inc.
10 Moulton Street
Cambridge, MA 02174
USA
Work Phone: +1 617 873-3087
Work Fax: +1 617 873-5620
Email: ccarroll@nnsc.nsf.net

Ms. April Marine
Research Associate
Work Address:
SRI International
NISC, Room EJ294
333 Ravenswood Avenue
Menlo Park, CA 94025
Work Phone: 1 415-859-5318
Work Fax: 1 415-859-6028
Email: april@nisc.sri.com

Internet Service Providers

ANS
Joel Maloff
Vice President - Client Services
Advanced Network and Services
2901 Hubbard Rd.
Ann Arbor, MI 48105
Work Phone: +1 313 663 7610
Email: maloff@nis.ans.net

BARRNET
William Yundt
Pine Hall Rm. 115
Stanford, CA 94305-4122
Work Phone: +1 415 723 3104
Email:
gd.why@forsythe.stanford.edu
Work Fax: +1 415 723 0010

CERFnet
Susan Estrada
San Diego Supercomputer Center
P.O. Box 85608
San Diego, CA 92186-9784
Work Phone: +1 619 534 5067
Email: estradas@sdsc.edu
Work Fax: +1 619 534 5167

CICnet
Michael Staman
President
ITI Building
2901 Hubbard Drive Pod G
Ann Arbor, MI 48105
Email: staman@cic.net
Work Phone: +1 313 998 6101
Work Fax: +1 313 998 6105

Colorado Supernet
Ken Harmon
CSM Computing Center
Colorado School Mines
1500 Illinois
Golden, Colorado 80401
Work Phone: +1 303 273 3471
Email: kharmon@csn.org
Work Fax: +1 303 273 3475

CONCERT (Communications for NC Education, Research, and Technology)
Joe Ragland
P.O. Box 12889
3021 Cornwallis Road
Research Triangle Park, NC 27709
Work Phone: +1 919 248 1404
Email: jrr@concert.net
Work Fax: +1 919 248 1405

Cornell University
Martyne Hallgren
Room 626
Engineering and Theory Center Bldg
Cornell University
Ithaca, NY 14853-3801
Work Phone: +1 607 254 8767
Email:
matyne@mitchell.cit.cornell.edu
Work Fax: +1 607 255 0823

CREN
Jim Conklin
EDUCOM
1112 16th Street NW
Washington D.C. 20036
Work Phone: +1 202 872 4200
Email: conklin@bitnic.bitnet
Work Fax: +1 202 872 4318

CSUNET
Chris Taylor
Manager, Network Technology
Office of the Chancellor
Information Resources and Technology P.O. Box 3842
Seal Beach, CA 90740-7842
Work Phone: +1 213 985 9669
Email: chris@calstate.edu
Work Fax: +1 213 985 9400

DataNet
Seppo Noppari
Telecom Finland
P.O. Box 228
Rautatienkata 10
33101 TAMPERE
Finland
Work Phone: +358 31 243 2242
Work Fax: +358 31 243 2211
Email: Seppo.Noppari@tele.fi

JVNCnet
Sergio Heker
6 von Neumann Hall
Princeton University
Princeton, NJ 08544
Work Phome: +1 609 258 2411
Email: heker@jvnc.net
Work Fax: +1 609 258 2424

Merit
Eric Aupperle
Merit Network
2200 Bonisteel Blvd.
Ann Arbor, MI 48109-2112
Work Phone: +1 313 764 9423
Email: ema@merit.edu
Work Fax: +1 313 747 3745

MIDnet
Dale Finkelson
29 WSEC
University of Nebraska
Lincoln, NE 68588
Work Phone: +1 402 472 5032
Email: dmf@westie.unl.edu
Work Fax: +1 402 472 5280

MRNET
Dennis Fazio
Executive Director
The Minnesota Regional

Network
511 11th Avenue South,
Box 212
Minneapolis, Minnesota 55415
Work Phone: +1 612 342 2570
Email: dfazio@MR.NET
Work Fax: +1 612 344 1716

NCAR
Joseph H. Choy
P.O. Box 3000
Boulder, CO 80307-3000
Work Phone: +1 303 497 1222
Email: choy@ncar.ucar.edu
Work Fax: (303) 497-1137

NEARnet
John Rugo
Accounts Manager
BBN Systems and Technologies
10 Moulton Street
Cambridge, MA 02138
Work Phone: +1 617 873 2935
Email: jrugo@nic.near.net

NET ILLINOIS
Ed Krol
University of Illinois
Computing Services Office
1304 W. Springfield
Urbana, IL 61801
Work Phone: +1 217 333 7886
Email: e-krol@uiuc.edu

NevadaNet
Jim Williams
University of Nevada System
Computing Services
4505 Maryland Pkwy
Las Vegas, NV 89154
Work Phone: +1 702 739 3557
Email: greyfox@nevada.edu
Work Fax: +1 702 739 3791

NorthWestNet
Eric S. Hood
Executive Director
NorthWestNet
2435 233rd Place NE
Redmond, WA 98053
Work Phone: +1 206 562 3000
Email: ehood@nwnet.net

NYSERnet
Jim Luckett
NYSERNET INC
111 College Place
Room 3-211
Syracuse, New York 13244
Work Phone: +1 315 443 4120
Email: luckett@nysernet.org
Work Fax: +1 315 425 7518

OARnet
Alison A. Brown
Ohio Supercomputer Center
1224 Kinnear Road
Columbus, Ohio 43085
Work Phone: +1 614 292 9248
Email: alison@osc.edu
Work Fax: +1 614 292 7168

Onet
Eugene Siciunas
4 Bancroft Ave., Rm. 116
University of Toronto
Toronto
Ontario M5S 1A1
Canada
Work Phone: +1 416 978 5058
Email:
eugene@vm.utcs.utoronto.ca
Work Fax: +1 416 978 6620

PREPnet
Thomas W. Bajzek
530 North Neville Street
Pittsburgh, PA 15213
Work Phone: +1 412 268 7870
Email: twb+@andrew.cmu.edu
Work Fax: +1 412 268 7875

PSCnet
Eugene F. Hastings, II Pittsburgh
Supercomputing Center
4400 5th Avenue
Pittsburgh, PA 15213
Work Phone: +1 412 268 4960
Email: hastings@psc.edu
Work Fax: +1 412 268 5832

PSINet
William L. Schrader
President & CEO
11800 Sunrise Valley Drive
Suite 1100
Reston, VA 22091
Work Phone: +1 703 620 6651
Email: wls@psi.com
Work Fax: +1 703 620 4586

SDSCnet
E. Paul Love, Jr.
San Diego Supercomputer Center
P.O. Box 85608
San Diego, CA 92186-9784
Work Phone: +1 619 534 5043
Email: loveep@sdsc.edu
Work Fax: +1 619 514 5152

Sesquinet
Farrell Gerbode
Office of Networking and
Computing Systems
Rice University
Houston, TX 77251-1892
Work Phone: +1 713 527 4988
Email: farrell@rice.edu
Work Fax: +1 713 527 6099

Sprint International
Bob Doyle
SprintLink Product Manager
Sprint International
13221 Woodland Park Drive
Herndon, VA 22071
Work Phone: +1 703 904 2167
Work Fax: +1 703 904 2680
Email: rdoyle@icm1.icp.net

SURAnet
Jack Hahn
1353 Computer Science Center
University of Maryland
College Park,
Maryland 20742-2411
Work Phone: +1 301 454 5434
Email: hahn@umd5.umd.edu

SwipNet
Olle Wallner
SwipNet AB
P.O.Box 6048
S-164 06 KISTA
Sweden
Work Phone: +46 8 924040
Work Phone (direct): +46 8 924058
Work Fax: +46 8-963695 Email:
wallner@swip.net

THEnet
Tracy LaQuey Parker
Computation Center
University of Texas
Austin, TX 78712
Work Phone: +1 512 471 5046
Email: tracy@utexas.edu

UUNET Communications Services
3110 Fairview Park Drive
Suite 570
P.O. Box 2324
Falls Church, Va 22042
Work Phone: +1 703 876 5050
Work Fax: +1 703 876 5059
Email: uunet-request@uunet.uu.net
Email (uucp): uunet!uunet-request

VERnet
James A. Jokl
VERnet
Academic Computing Center
Gilmer Hall
University of Virginia
Charlottesville, VA 22903
Email: jaj@virginia.edu
Work Phone: +1 804 924 0616
Work Fax: +1 804 982 4715

Westnet
Pat Burns
UCC
601 S. Howes, 6th Floor South
Colorado State University
Fort Collins, CO 80523
Work Phone: +1 303 491 7260
Email:
pburns@yuma.ACNS.ColoState.EDU
Work Fax: +1 303 491 2293

Some Thoughts on NSFnet Router Design

by Dr. W. I. Nowicki

I was not privy to the NSFnet details, but have some relevant opinions. Disclaimer: I am not unbiased, since our company develops products for the RS/6000 (and many others) and I am a token IBM stock-holder. Feel free to use in your newsletter, as long as you keep a disclaimer that these are all personal opinions:

My last job was doing network software for Sun, where I fought (but lost) to get Sun into the router market. Actually, Sun is probably the third or fourth largest router vendor, but does not "market" or "develop" them. I like to distinguish between using standard and special purpose platforms (I used to call these "open" and "proprietary" but those words display my bias too much), and using a dedicated or shared box. All four combinations make sense, and have been used at one time or another. For example, the boxes Cisco makes are special purpose, and dedicated. At one time they sold boxes that could be routers and terminal servers, for example, at the same time, and those would be special purpose but shared. If you take a file server and slide in a second Ethernet interface, or use point-to-point software for a wide area network link, then you have a classical general purpose shared router.

Unfortunately these general purpose shared routers got a bad name in the early days when people tried to use VAX-11/750s as timesharing machines and routers, when they had the horsepower to do neither. But now with commodity hardware they make sense in many situations, and Sun and DEC have quietly sold tens of thousands of them. What IBM has done for NSFnet are general purpose but dedicated routers.

	Platform: General Purpose	Platform: Special Purpose
Shared	Sun, etc.	Cisco Terminal Server/Router combination
Dedicated	NSFnet NSS	Cisco Proteon, Telebit, etc.

The technical reason to use a special purpose platform is that one can optimize for routing performance. However, the technical reason to use a general purpose platform is that one "rides the curve" of commodity hardware; that is, you more easily take advantage of the cost / performance improvements that can be amortized over a much larger volume. For example, the SPARCstation-2 is a low-cost commodity machine that makes a fine router. And once the next SPARCstation model comes out, the software investment is protected with very little new work. Low-cost routers can also be easily made from 386 PCs.

Primarily, it comes down to a market positioning issue: IBM is clearly in the general purpose platform business, while Cisco ONLY wants to sell special purpose boxes (with higher margins). DEC is an example of the effect of commodity hardware pricing: it is turning into a service business, so the kind of platform does not really matter. Sun evidently wants to sell high-volume desktop boxes, so is not interested in the router market as a separate entity.

Right now the customer is U.S. Congress, who wants photo-ops. If NSF funded individual researchers or institutions who would buy Internet service on the open market, the customers would then be those researchers and institutions. Given the true cost of the phone lines, few customers could afford 45 Mbps service. IBM would then concentrate more on scalable architectures for large numbers of low and medium bandwidth connections, for which using RS/6000 based routers makes even more sense, once the POWER-PC chips come out to make them true commodity CPUs. I think this future is inevitable, it is just a question of time.

Another important advantage of using a general purpose platform is that new administration tools do not need to be used. Of course, companies like Cisco and DEC can make some extra money by selling the special purpose platforms, and then charging extra for training and service, while IBM usually bundles service into their prices, so prefer uniform administration. It is also possible to use general purpose CPUs and operating systems, a standard control bus, but some custom hardware at the lowest level for better performance. This combines the advantages of both approaches, and seems to be what IBM is doing with its '960 interface boards.

In conclusion, IBM is probably doing the right thing in my opinion The only problem is with the politics of the situation. The funding comes through a "cooperative agreement" filtered down through many levels
.
There is also confusion between this being a research prototype or a production network. My hope is that in the future a more truly open competitive market would force them to be more responsive to the "customer". Right now the customer is U.S. Congress, who wants photo-ops. If NSF funded individual researchers or institutions who would buy Internet service on the open market, the customers would then be those researchers and institutions. Given the true cost of the phone lines, few customers could afford 45 Mbps service. IBM would then concentrate more on scalable architectures for large numbers of low and medium bandwidth connections, for which using RS/6000 based routers makes even more sense, once the POWER-PC chips come out to make them true commodity CPUs. I think this future is inevitable, it is just a question of time.

ANS & 501(c)(3) (Cont. from p. 1.)

2. They had been denied status. He asserted that six months was quite long enough for a normal application to be processed, adding that it was certainly not a case of whether or not the original application had been filed correctly. There had been plenty of time to rectify that. But he was also careful to assert that there were two other possibilities.

3. Given the nature of the players in this case and the controversy surrounding it, the IRS might be going over its decision process very slowly and meticulously.

4. The IRS might also be negotiating some details such as the effective date or type of status with ANS.

5. A final possibility might be a combination of 3 and 4.

While he asserted that in a normal case six months was plenty of time for an IRS decision, other sources have disagreed pointing out that in their experience 18 months for an IRS ruling can be quite frequent.

So this part of the ANS saga may have quite a few months more to run. In any case the COOK Report will be checking back with the IRS and updating the ANS status -- probably monthly.

While the IRS is figuring out what to do, perhaps it should be reviewing the ANS cooperative and gateway agreements. The August issue of *Data Communications* has an interesting article on the internet from the European perspective. After saying that OARnet has a gateway agreement with ANS CO+RE, it goes on to say on page 107:

"Charges for connecting other network providers to the ANS backbone vary according to access speed and usage type. For instance a T-1 connection costs $43,000 a year for non commercial traffic and 47,000 a year for a mixture of commercial and non commercial traffic. Mid level service providers pay ANS an additional annual fee for each connection to their network; a 56-kbit/s connection costs $2,000 a year for non-commercial traffic and $2500 a year for a mixture of commercial and non commercial traffic. ANS puts the extra fees for commercial traffic in a fund that is used to subsidize Internet developments. It keeps the base fee however."

From this I infer that OARnet which didn't previously have a direction to the backbone is paying 40 some odd thousand extra per year for a direct connection. In this case I believe through Cleveland. I wonder whether the only mid-level service providers who would consider this hefty extra annual fee are those roughly 16 who do not have direct connects to the NSFnet backbone?

Question 1: does money go to the infrastructure pool only if OARnet buys the commercial connect? $47 instead of 43 thousand. $4000 to the infrastructure pool? $43 thousand for ANS?

Questions 2-8 follow from the statement: "mid level service providers pay ANS an extra fee for each connection to their network." 2,000 to 2500 dollars extra for each 56 kbit line.

Question 2: *Which* mid-levels? All 32 or is it 33 by now? Only those mid-levels that run a backbone node? Or only those mid-levels that sign a gateway agreement with ANS. Or only those mid-levels like OARnet that take a direct connection to ANS?

Question 3: do mid-levels with T-1 connects pay a $4000 surcharge for each connect? Or only $2000? Note that many mid-levels have 100 or more institutions connected. So we are talking real money at an average of $2000 per institution. Is money paid to the infrastructure pool *only* when a commercial hookup is tallied? And is the money paid only the difference between the commercial and non commercial rate ($500)?

Question 4: Is one to assume that the cost to OARnet for its ANS hookup is the 40 odd thousand *plus* an average of $2,000 per connected institution? That would make the direct connect very expensive indeed.

Question 5: Do mid-levels like NYSERnet that have gateway agreement with ANS have to pay the extra money per for each institution on their network?

Question 6: The charges are $2000 for non commercial traffic and $500 higher for commercial.and non commercial mixed. I can see why a mid-level might pay the mixed commercial & non commercial rate for the right of its client to be free of backbone AUP. *Why* would a mid-level consent to a $2000 surcharge for each non commercial client?

Question 7: What networks are actually being affected? Someone said that half the mid-levels had signed the ANS agreements. Are those that have not still blocked from accessing commercial ANS traffic?

Question 8: If I were operating such a mid-level, why would I sign up and ask for commercial status for some of my connected institutions, if it meant that they would be blocked from the institutions connected to the one half of the mid-levels that have NOT signed the ANS agreements? After all we have seen that virtually any organization that uses a mid-level can also send packets across the backbone just so long as it takes a scout's honor's pledge that it's packets support research and education.

Question 9: The famous infrastructure pool: who is responsible and how much is it worth? From what I had read when it was first discussed about a year ago, the NSF appeared to have some oversight responsibility. ANS is a would-be non profit organization. As part of their keeping of the public trust, should they not report to the community the value of the pool now more than a full year after its inception?

Question 10: What should ANS, in return for its would be non profit status return to the community? What *is* it bringing to the table? Early speeches gave the impression that IBM and MCI would be using it to create (implied donate) a network infrastructure for American research and education. Therefor it would be eligible for special status that its competitors, ANS, PSI, Sprint, and the mid-levels don't have.

Has ANS done anything to *earn* its would be non profit status? Or has the NSF merely vested it with the right to levy surcharges, which in the end will have to be paid through higher connect prices to the network, but in the meantime appear to be the principal justification for would be 501(c)(3) status because they are being placed in an "infrastructure pool?"

The *COOK Report* really would like answers to these questions: from ANS, from the NSF, from the mid-levels, or from members of the network community.

Subscription Information

Twelve "hardcopy" issues of *The COOK Report* will cost $75.00 for individuals and companies or government agencies with operating revenues of $10 million or less per year. The cost for larger entities is $300.00. *The COOK Report* will be available in electronic and hardcopy form to subscribers at the corporate level for $375 (site license). These subscribers may duplicate hardcopy or electronic for distribution within their company only. Subscribers at the $300 level and above will receive for no extra charge the 35 page (170,000 characters) desk top published report: *The National Research and Education Network: Whom Shall It Serve?* Others may order this document separately at $25. Foundations and non-profit corporations may subscribe for $150 a year.

Foreign hardcopy subscriptions are $125. Foreign corporations pay the American rate plus $50 a year extra for postage and handling. The *COOK Report* is sold by the year. Current subscribers get all issues of year one.

Send payment by check (payable to Gordon Cook) at the following address.

Gordon Cook, President
COOK Network Consultants
431 Greenway Ave
Ewing, NJ 08618, USA

Telephone & fax (609) 882-2572
Internet: cook@tmn.com

The Cook Report on Internet - NREN
COOK Network Consultants
431 Greenway Ave.
Ewing, NJ 08618