The COOK Report on Internet → NREN

Vol. I, Number 5, August 1992

The Network as Testbed for Router Development

Divining the Meaning of Backbone Solicitation in the Context of the T-3 Router Woes

During mid-summer the network was marked by three developments: attempts to divine what the NSF backbone solicitation draft really means; the publication on July 7th in *Communications Daily* of the first results of an investigation into the relationship between the NSF, MERIT, and ANS-IBM; and the growing awareness of IBM's problems in trying to get meaningful results from its use of the NSFnet as a testbed in router development.

The Solicitation

Awareness grew that the solicitation draft released early in June is nothing more than an outline of a "philosophy" for how the network should be organized. The gaps that must be filled in before the major players can react are extensive. People were heard to say that, based on the information in the solicitation, some of the potential bidders were having difficulty in understanding well enough what was wanted in order to write serious critiques. At mid month it was reported that the NSF had received no responses. (August 3rd is the response deadline.)

The *Communications Daily* article outlining the troubles of the IBM RS 6000 T-3 router emphasized further the weakness of the cost sharing arguments. Policy makers should ponder the current network experience and ask themselves whether the driving purpose of the network is really to increase US technology competitiveness. If the answer is yes, they might note that a company like IBM, that has seldom been a technology leader, would have ample reason to invest 3 or 4 corporate dollars for every federal dollar received in order to play technology catch up with the rest of the market place. (Estimates are that the network and routing authority for which the NSF will pay $10 million a year will cost the providers $40 million a year to deliver.)

Meanwhile a market leader like a Cisco would find considerably less reason to pay for the privilege of developing technology that it was doing on its own successfully anyway. If the use of federal money is really intended to push the leading edge of the technology and to help researchers by delivering a faster network, the procurement process must be changed.

Communications Daily: IBM Deploys "Crippled" Technology

On July 7th *Communications Daily* published a lengthy article on the shortcomings of the National Science Foundation's efforts to upgrade the NSFnet backbone to T-3 speeds. The article's major conclusions were that: "(1) Network upgrade funded with $28-million govt. grant, never has materialized fully. (2) Proprietary IBM technology developed for network was deployed before being tested. (3) Contradictory statements of officials responsible for overseeing network as

(Cont. on p. 2)

The Community Learning and Information Network Revisited

by Dave Hughes

Recently for *The Cook Report on Internet -> NREN*, I reviewed a draft concept plan for the Community Learning Network, CLN, as it was presented to me by one of its early advocates. What I read was not impressive in its details of proposed implementation. But some of its strategic concepts intrigued me, for they seemed to be the first attempt at a solution to a series of general education and training problems that were being approached on a national basis with a feasible funding concept.

So I did not dismiss it out of hand, as some have done, and continue to do.

Since then, there have been a number of developments and changes to the

(Cont. on p.9)

On the Inside

The gaps [in the solicitation] that must be filled in before the major players can react are extensive.

to its overall performance were made in forums such as congressional hearings, public conferences, technical symposia. (4) govt. grant establishing network included no technical specifications, checks or balances, leaving granting agency, NSF, without any control over network it's paying for."

According to interviews with router design engineers who protested to management that "they couldn't possibly produce full functioning T-3 router without extensive research and testing," current ANS President Al Weis, "while at IBM, spearheaded router's design and was person who promised NSF that IBM could deliver stable T-3 router 'within months'." However Weis informed Brock Meeks, the author of the article that "IBM's commitment (to develop router) was not made by person on my level. It was made at very high levels within IBM."

According to interviews with router design engineers who protested to management that "they couldn't possibly produce full functioning T-3 router without extensive research and testing," current ANS President Al Weis, "while at IBM, spearheaded router's design and was person who promised NSF that IBM could deliver stable T-3 router 'within months'."

Routers: Some of the Stakes for IBM

Not being privy to IBM's inner most strategic decisions, it is impossible to surmise with certainty the reasons for the decisions made by the company in the development of its routers. During the last decade and perhaps longer, the company has seldom been the first to market with a new technology. It has adapted a pattern of waiting until a new technology trend becomes clear and then swinging its massive size and marketing ability behind the trend to take over and dominate the technology. Such a strategy was fabulously successful with the first generation of personal computers between 1981 and 1984.

However, as the speed of industry development increases, there are signs that IBM may have waited too long to begin serious multi-protocol router development. As we have reported before, some Fortune 1000 managers are eliminating SNA networks entirely. The trend is for them to fold DECnet and SNA nets into a single corporate network using multi protocol routers brought to the marketplace by companies such as Cisco, Proteon, and Wellfleet in the last two years. The technology solutions of non IBM companies have made it possible for those with IBM mainframes using front end controllers that may cost from 80,000 to 300,000 dollars to replace them with multi protocol routers costing a tenth as much.

The Legislative Context of Agency Rivalries

In early 1990 when IBM and the NSF began their discussions that led to the May 1990 approval by the National Science Board of the decision to upgrade the NSFnet backbone, these trends may not have been as clear as they were two years later. Nevertheless there was political pressure to hurry because on Capitol Hill those backing S1976, a bill defining the duties of the Department of Energy had, during the process of developing a committee report, changed the bill's language from ordering the Department of Energy to cooperate with the NSF in the establishment of an NREN to ordering it to establish the network directly.

In view of what has happened since then, the language of the report appears to have been an accurate predictor of future events. Noting that OSTP wanted FRICC (the predecessor of the Federal Networking Council) to develop to coordinate the implementation of NREN, page 13 of the committee Report on S1976 stated that "FRICC can discuss these issues, but as an informal inter-agency committee, it lacks the authority to implement decisions. . . . Even if FRICC has the authority to create the NREN, the future of the NREN will be uncertain."

"Since FRICC is an inter-agency committee, no one agency is solely responsible. Each agency will vie for its own individual agendas through FRICC. No one agency has the authority or responsibility to resolve the inter-agency conflicts and demands. The tough choices that have to be made will more likely be deferred than resolved."

"A single lead agency with sole responsibility is needed to take charge of the NREN. . . . The success of a national high performance computing initiative also depends ultimately on the management structure placed up it and not on the spending levels." In view of the fact that it took the FNC two years to put together an inter-agency management plan for federal network coordination having nothing directly to do with NREN, and in view of the tortuous difficulties in rebiding its own backbone that the NSF has become tied up in, the words of the committee report completed on June 27, 1990 appear to have been very wise.

[Said the committee on behalf of the Department of Energy:] "A single lead agency with sole responsibility is needed to take charge of the NREN. . . . The success of a national high performance computing initiative also depends ultimately on the management structure placed up it and not on the spending levels."

Meanwhile what was now opposing legislation (S272) stated that the National Science Foundation "in cooperation with" NASA, DOE, and DARPA "shall provide for the establishment" of a multi-gigabit per second NREN. The next year was spent in jockeying within the Senate designed to get the backers of S343 (the 102nd Congress version of S1976) to be willing to fold their demands on behalf of the Department of Energy's control over NREN into the vision of S272 that gave the role to the NSF. The arguments posed to counter those favoring DOE on the Hill were that, because the NSFnet was larger, it would reach a broader segment of the nation's research and education population. Besides it was now faster since the NSF announced in June 1990 that it was upgrading 8 nodes to T-3 immediately. This was followed with a September announcement that three new nodes and the remaining five old ones would be upgraded to T-3.

In May of 1990 when the Science Board approved the upgrade, IBM

knew that it had to deliver the routers and do it in a hurry. According to the interviews of involved IBM design engineers conducted by Brock Meeks for his *Communications Daily* article, a "SWAT team" of engineers was detailed to attack the project. One is certainly tempted to suppose that, for political reasons, speed was essential.

Design Choices for the RS/6000

In developing routers for the T-1 network IBM had used its RT workstation. Faced with a tight, "as soon as possible," deadline for the T-3 NSFnet, it choose to do the same thing with the RS6000 workstation. In October of 1991 *Networking Management* was critical of the decision. In a news analysis piece titled "IBM attempts router market scare with time bomb" it stated:

"Although the RS/6000 incorporates RISC in an elegant and powerful package, no work station platform has out-performed router-specific architectures for internetworking applications -- especially now that these dedicated routers are already employing RISC. Modifying a workstation designed for graphics and number crunching to specifically route traffic across a multi-vendor network is an architectural stretch."

"Modifying a workstation designed for graphics and number crunching to specifically route traffic across a multi-vendor network is an architectural stretch."

What was IBM doing with the 6000? By the fall of 1991 it became clear that it was using it as a design foundation for the commercial multi-protocol router announced in August and now known as the 6611. The same *Network Management* article stated: "IBM's recent teleconference with trade journalists about its routers was a desperate ploy by Big Blue to stem the tide of dollars to successful internetworking vendors such as Cisco, Wellfleet, 3Com, ACC, Vitalink and Proteon. Multi-protocol routers are data communications linchpins in the downsizing trend from mainframes to heterogeneous distributed computing systems built with muscular workstations."

IBM announced that a delivery date for the commercial version of the RS/6000 would be set during the winter. By January of 1992 the delivery date for the 6611 had been set for late June. Meanwhile we wondered why IBM would take such an architectural stretch given the now apparent performance problems of the RS/6000.

A question to the com-priv list brought some enlightening responses. According to Karl Denninger - author of the AKCS conferencing software - You can achieve high performance with good results in dedicated architectures by switching packets on a multi-port interface board without ever touching the backplane. On a workstation platform, this is much harder to do. "However, on a workstation, assuming you have the routing code (gated, etc) already working, you only have to write a device driver. Not that bad of a job. The rest is already done. This saves 'mongo' time."

"The price? Performance and reliability. Booting from a fixed disk, and having to require one in the machine, is in my opinion a serious reliability problem - enough for me to eliminate the equipment from consideration. The Mean Time To Repair goes from a few minutes to a few hours when you do this, and the Mean Time Between Failures drops by at least half! Booting from floppy is better (backing that up and having a spare is easy), and from Programable Read Only Memory with TFTP capability is best."

"The price? Performance and reliability. Booting from a fixed disk, and having to require one in the machine, is in my opinion a serious reliability problem - enough for me to eliminate the equipment from consideration.

Denninger added that if a router booted from PROM, it was nearly impossible to be left unable to restart unless the entire router is destroyed). In such a case with quick-swap cards you have a router that, given spare parts on-site, can be brought back online in 5 minutes or less in the event of failure. "A fixed-disk box requires that I have preloaded and preconfigured spare machines laying around depreciating instead of boards, and drives up the total ownership cost. Ugh."

"Swapping Cisco boards is a few-minute operation. The only exceptions in the AGS+ are the large 'hair dryer' fan and the power supply. Neither of those take more than a half-hour worst-case; the cards are near-instant. I've never had a fan or power supply fail in these boxes to date. Cards, yes. How long does it take to restart a 6611 if the internal disk drive fails? Can you be without the network link for that long?"

"By the way," Denninger concluded, "I did an evaluation process for a national TCP/IP backbone about 6 months ago (private network; 5,000 distinct sites planned at full expansion). The first conclusion that fell out of the MTBF/MTTR analysis is that you want *no* spinning magnetic media in the field."

Stan Hanks, Principal Scientist, Technology Transfer Institute, added the following insight into the simplified software development that is available to those taking the workstation rather than the specialized architecture option.

"Question: if you need to add a new type of interface, which is easier? Answer: the UNIX-based method. Why? Because you get a simpler methodology with which to develop the critical device driver. And the rest of the software can't be impacted by the new driver." The hoped for dividend is early market entry since it is "much easier to add a new device driver to a unix system than to craft a whole new set of system software from scratch.

We surmise that counting on the earlier and easier technology solution was the only way that could give IBM any hope of delivering a high speed router to the new backbone 'within months' as promised by Al Weis. Unfortunately for all concerned the months may be turning into years. The RS/6000 router development project has been underway for more than two years. Although the routers are within the past two months functioning with fair reliability for the first time, they are still very slow. Average speed on the T-3 "45 megabit" backbone appears to be in the 5 to 8 megabit range. A yet-to-be-resolved problem with FDDI interface cards apparently prevents two FDDI local area networks from being connected via the backbone at greater

that ethernet (10 mgbs) speeds. We have recently learned that the RS/6000 routers are currently, and purposefully, engineered to run at 22.5 mbps. They can, IBM claims, be "re-engineered" on an "overnight" basis to run at full, clear channel, T3. But they "don't have the need for it right now, [because] the network isn't anywhere near pushing the T3 links that fast."

In front of academic audiences and Congress ANS has been very clear that its over riding purpose is bring high speed networking to the research and education community. With the T-3 network, it has an opportunity to do this. Unfortunately it is not delivering. It has also asserted its independence from IBM. One would think that a good way for it prove its sincerity would be to remove the RS/6000 boxes and install Cisco AGS+ routers. Challenged on this point on the com-priv list, it has remained silent. It's silence and apparent inaction seem to point to its ties to IBM's networking testbed needs as the principal reason for its continued existence.

The Saga of the 6611

Meanwhile just days before the promised end of June ship date, IBM announced a 90 day postponement in the availability of its 6611 commercial multi-protocol router. If one looks carefully at a lengthy article in the March 1992 issue of *Data Communications* outlining the router's capability, one can begin to understand what some of the reasons may be.

The 6611 is being called upon to do two jobs. Its first market is the interconnection of LANs over wide area networks. For this task it handles TCP/IP, and Novell's IPX and SPX while using proprietary means to handle SNA, Netbios and LAN broadcast messages. Because of this, its interoperability with the routers of other vendors may not be extensive. Meanwhile its second market is IBM's huge installed SNA base. It is thought that the majority of these sites may be moving to LAN and SNA network interconnects. The demands of these two markets are not very compatible.

The 6611 is being called upon to do two jobs. Its first market is the interconnection of LANs over wide area networks. . . . . Meanwhile its second market is IBM's huge installed SNA base. It is thought that the majority of these sites may be moving to LAN and SNA network interconnects. The demands of these two markets are not very compatible.

While the 6611 attempts to solve some of the problems (such as excess overhead and timer problems) created by SNA in router based networks, the apparent need to make it fit within the SNA world and not replace completely IBM's front end processors appears to have led to design compromises. A v.35 interface needed to handle dial up 64 kbs traffic is not available. This limits its usefulness in many LAN interconnect schemes. "The 6611's routing software comes preloaded on an internal hard disk. The use of the hard disk poses a reliability problem: if the disk fails the router fails. IBM is well aware of this but it isn't saying when it will address the problem."

As do competing routers from Cisco and Wellfleet, the 6611 routes otherwise unroutable SNA by encapsulating it inside TCP/IP packets. According to Robin Layland, an engineering consultant in charge of network architecture at the Travellers Insurance company, "the 6611 addresses problems of Logical Link Control (LLC) overhead, timers, router network flow control, source routing hop counts and broadcast messages. No other router addresses all these problems."

The router handles the problem of LLC overhead by splitting the session into three parts and not allowing the LLC data to traverse the wide area part of the network. However, "splitting the session into three parts also complicates the routing process. The router must keep track of what each part is doing and coordinate all three processes." One must assume that this balancing act places an extra load on the router and may help to account for its apparent slowness when matched against those of competing vendors.

Apparently the session splitting was necessary to handle another vexing problem of the IBM hardware software world. Token ring LANs are predominant in IBM shops, However, "as defined by IBM, token ring source routing only allows seven hops between two workstations. In other words only eight LANs counting the original ring can be connected. With LAN internetworks starting to span the country and the world, a seven hop maximum limits network expansion." Breaking the LLC session into three parts provides a work around to the limits on source routing hop counts.

Unfortunately the hop count work around comes with a down side. In the event of network problems, a single LAN analyzer is unable to see the entire network. Two analyzers and a management system are necessary to troubleshoot the network -- an unfortunate problem since the 6611 lacks "even basic problem determination tools."

The 6611 also handles SNA traffic based on SDLC (synchronous data link control) "by breaking the SDLC session into three parts." In general SDLC overhead is handled at the router rather than being sent over the network.

It doesn't take a router design engineer to surmise that the network overhead reduction features just discussed -- and they are only a portion of those in the full article -- will put a load on the throughput capacity of the router itself. According to Layland: "all those overhead reduction features will deliver real benefits, but it remains to be seen how well the 6611 will perform in the quarter mile. Cisco, Wellfleet and others already sell devices with very high packet rates and IBM is well aware of the stiff competition in this area." Leyland concludes: "with its 6611 introduction in late January, IBM pulled off an impressive debut. Now it faces two even more difficult tasks: delivering all the features it promised and gaining share in an intensely competitive market. IBM delivery of the 6611, scheduled for June 26th, will offer the first indication of whether it is up to those challenges."

The NSFnet Testbed Hasn't Helped

IBM didn't make its delivery date, announcing on June 26 a three month delay. The page one major story in *Computerworld* was headlined: "IBM delays new router, loses face. More testing needed, few users waiting." *Computerworld* added: "IBM bit off more than it could chew in attempting to turn around a complex initial offering on an 18 month development cycle

according to users and analysts."

Robin Leyland in the *Data Communications* article from which we have extensively cited offered some further insight: "Realizing it had to work fast, IBM sought the help of established router vendors, most notably Wellfleet. When that effort failed IBM instead turned to the "hot box" it had already developed in house -- the RISCX based 6000 workstation." But what may not have been evident at during the February time frame in which the Leyland article was prepared, is very apparent now -- the RS/6000 is not a "hot box".

IBM - behind the market again - was forced to play catch up. Whether from a lack of understanding of the TCP/IP LAN interconnect market, or from the lack of an adequate engineering base because of its recent significant downsizing, it tried to make a single product play catch up in two basically incompatible areas: LAN interconnects, and SNA.

IBM has recognized for at least two years that it is faced with the need to deal with a major technology shift from proprietary network protocols and large mainframe computers to a world of open standards and distributed computing. Unfortunately for Congressional strategists, it's possession of a national testbed in the form of the NSFnet doesn't seem to have affected the outcome -- which appears to be that of a huge company caught in a web of contradictions with which it can no longer adequately cope.

Sprint Announces Nationwide Commercial TCP/IP Internet Service

On July 22 US Sprint Announced the availability of Sprintlink offering commercial links to the Internet. The new service grew out of Sprint's work in managing a network to link the National Science Foundation's NSFnet with the-French scientific and research network, INRIA, and the Scandinavian research network, NORDUNet, under TCP/IP.

The rise of SprintLink will put additional pressure on the Internet to solve a growing shortage of TCP/IP addresses. (In June the Internet Architecture Board made what turned out to be an abortive attempt to adopt an OSI addressing protocol.) Sprint's plans for SprintLink are certain to increase traffic on the network still further. The speed of SprintLink is 9600 baud, In the United States Sprintlink is accessible by means of dedicated local loops from all of Sprints 270 points of presence.

IBM Successfully Opens Company Wide Internet Gateway

As many large corporations join the NSFnet, we have had relatively little information on corporate use, which until recently has tended to be limited to a small subset of employees at each company -- often because of concerns about the leakage of proprietary information outside the company. We suspect that corporations are increasingly giving all their employees access to the network. The *COOK Report* thanks Howard L. Funk of IBM for sending the material on which the following summary of IBM's experience with unrestricted employee access is based.

In May of 1990 IBM opened to all its employees a gateway between VNET its corporate network and the Internet. By March 1 1992 13,000 users within IBM were sending 240,000 messages per month through the gateway. The gateway was accomodating new users at the rate of 500 per month. The goal of the gateway is to "provide a mechanism to enhance the technical vitality in IBM's development organization."

Twenty-seven percent of all PhDs within the company were using the gateway as were 10% of those with two or less years of service. Usage declines with years of service down to about 3% of those in the company with 10 or more years of service. Over eight percent of all engineers and programmers and a slightly smaller percentage of programming managers used the gateway. IBM's Marketing Division is the largest single organizational user of the gateway. The IBM account manager for a major government laboratory reported that the his customer said: "the best thing you did for us last year was opening an Internet gateway."

Connectivity to the NSFnet, originally provided via NYSERnet and PSI, is now provided via ANS. An automated server allows employees to register themselves to use the gateway. The server checks the results with the corporate directory, and if the check is successful, send the new user more information on how to use the gateway. The automated registration system saves the costs of a network admistrator.

For security reasons only electronic mail and mail lists are available to users. When a user sends a message to an internet address, the gateway reminds him that the adressee is outside of IBM, and that rules involving use of the gateway to not send proprietary information or information that would violate US export controls must be followed. If the user indicates that he still wishes to procede, the message is transmitted.

SOME USER COMMENTS

"Thanks. We're finally part of the real world.!"

"It is almost trouble free and has revolutionized my relations with the academic world. For example seven academics and I have written/edited a book in just over a month on the network."

"If this service didn't exist, it would be necessary to invent it. It's an excellent example of IBM's committment to techno-glasnost.'

"very useful in correspondence concerning university contracts."

"It impresses customers that IBM is finally coming around to open systems."

"very important in the execution of my current contract. It provides a link with my US Govt. customer and subcontract developer."

"indispensable tool for communicating with my customer as we negotiate . . . a large joint study project. Internet access is necessary for me to get my job done effectively."

"I am very hopeful that this new openess of communication will encourage further cooperation. Perestroika?"

Who's Pulling the NREN Strings? Musings of an Electronic Cynic

a guest editorial by Dave Hughes,
Old Colorado City Communications, Colorado Springs, CO 80904

Editor's Note: When the July 7th *Communications Daily* article by Brock Meeks was republished on the com-priv mail list Jack Rickards the publisher of Boardwatch magazine entered a series of postings in which he argued that Meeks and Gordon Cook were unfairly disparaging the NSF, ANS and IBM and that everyone should stand aside and give IBM motre time to finish the job of upgrading the backbone to T-3 speed.

Dave Hughes' response to com-priv:

Well I don't think the issue Meeks presents is whether the first routers by IBM didn't perform to expectations, or whether open bidding by others - such as Cisco - would have produced better ones. Rather its his revelations, that the whole 'T-3' affair has come out of what smacks of an murky alliance between three entities alloperating with extra agendas besides just doing Patriotic Red White and Blue pioneering into high speed networking. There was a hell ofa lot of institutional, not to speak of commercial, self-interest involved, which has made this more than a simple case of making technical mistakes at the leading edge. The issue raised is whether bad technical decisions were made because of influences that should not have been there.

Its pretty clear that the NSF wanted 45mbit T-3's any way they could get it. In a big hurry, that short circuited a lot of the procedures designed to prevent bad mistakes. What was the rush? Why was the decision to throw another $14 million at the T-3 project done in such haste? Trying to beat the Russians.?

Or is it only a coincidence that Louisiana Senator Johnston, in late 1989 introduced Senate Bill 343 that would have made the *Department of Energy*, instead of the NSF, the 'lead agency' for advanced networking. That turf war couldn't have anything to do with it, could it? And the percieved (by the NFS) 'need' to get high speed networks up before anyone else could - or that bill went through?

Agencies in Washington - right now - are all fighting over who is going to get the big post-cold-war bucks for advanced computing and networking? Noble motives, insuring the best decisions, Jack?

Why did IBM's management overrule their own engineers who hollared they couldn't deliver 45mb T-3s in the time ANS reassured NSF they could have it? What was the rush? Would the networking world have come to an end had they said to NSF "Hey this is leading edge stuff. We can't promise those speeds fast. All we can do is the best we can, better than T-1, and after all, we are IBM, remember?"

Afraid they would lose to competition that would bid in a more open process? Eager to corner a monopoly on the sale of routers to the monopoly of the backbone? And lock up future profits by that oldest 'private enterprise' ploy of all - non-competitive guarantees from government? ANS can protest all it wants that it is an independent non-profit organization just doing good for the network nation. But who supplied the routers?

Where is your customary 'conservatism' Jack ? Competition, risk taking, entrepenership, and the grand American way? Or, as a journalist, are you cynically saying government in bed with big business has always been the way things are done? And that it 'doesn't matter?' so long as something is getting done. Gee. I'm still idealistic enough to think that when a government agency, acting in the public interest, and without special agenda's of its own, manages an important national project, we have a chance of getting the best, not just the most expedient.

And what about that entity who was in between them, Merit? Michigan Merit? How did they get off the hook, being the actual entity doing business with NSF, and contracting with whomever they choose? And *I* get mighty cynical when I note that the President of the University of Michigan, was also the Chairman of the National Science Board that made that rushed decision to throw more money at the T-3 project. The same U of Mich, with Merit on campus and rather initmately tied to it. And which sits on the ANS board, that installs the routers.

Hey, this whole thing has sure been convoluted, Jack. Best way for the US Government to do business in the public interest?

Then finally, when Brock Meeks quotes NEARnet, and NEARnet comes onto com-priv and bends over backwards to say how satisfied they are with the T-3 service they are getting. Couldn't be another case of 'Hey we gotta get along with big blue no matter how bad they are so we better say that reporter didn't get it right. Reporters are always held in low esteem, right?' Which makes me take anything NEARnet says from this point on with more than a grain of salt.

One can get *very* cynical, given this litany. Which has very little to do with router speeds.

JvNCnet Goes Private and Profit Making

On August 1 JvNCnet was spun off from its Princeton University home as a commercial profit making network. Its assets and liabilities will belong to a corporation, Global EnterpriseServices Inc. formed by network director Sergio Heker. The CIX can expect a new member. Letters were mailed to JvNCnet members on July 23rd notifying them of the change.

It will be interesting to see how many other mid-levels follow suit over the next year. Relying indefinitely on NSF subsidies would appear to be a poor choice on the part of those that wish to continue in existence. We wish JvNCnet good luck.

New Gore Legislation to Create National Information Infrastructure Has Serious Shortcomings

On July 1, 1992 Senator Gore submitted to the Senate a new bill designed to build a national information infrastructure and coordinate federal High Performance Computing and Communications activities in support of this effort. The subject matter is indeed very important to the future economic competitiveness of this nation. However the bill suffers from a number of apparently unrecognized contradictions.

What is to be Fedral Role in Marketplace Applications?

The bill focuses upon marketplace applications. As the senator's press release explained: "To help increase U.S. competitiveness and create jobs for Americans, the technology developed by the High-Performance Computing Act must be moved from the laboratories into the marketplace where it can be used. This new bill will make sure this happens by developing practical applications for advanced technologies and high-performance computing." Gore concluded: "Through the legislation I introduced today, hospitals, universities, research centers, factories, schools and libraries will be connected to advanced computer networks, sharing data, increasing U.S. competitiveness and improving the quality of life for all." Worthy goals for sure.

Unfortunately such goals appear to ignore a few harsh realities. While it has long been the practice of the government to fund the development of new technologies, it has been recent practice not to fund marketplace applications. Of course the question of a government industrial policy is relevant here. Current Republican thinking is that help to precompetitive technology development is appropriate. A technology that can be applied in the marketplace is by definition competitive not precompetitive.

The picture is confusing. While on the one hand the current wisdom is that the government should not pick marketplace winners, the policy of "cost sharing" on the network backbone means that the government is giving the winner an advantage. There will be many who supported government assistance for HPCC R&D who will look at the follow on legislation and wonder why the government and not the marketplace is now called on to build infrastructure.

In this context the definitions given in the legislation for grand application and information infrastructure are interesting: "(1) The term 'Grand Application' means an application of high-performance computing and high-speed networking that will provide large economic and social benefits to a broad segment of the Nation's populace. (2) The term 'information infrastructure' means a network of communications systems and computer systems designed to exchange information among all citizens and residents of the United States."

Some will question the wisdom that government agencies can define a Grand Application better than the market place. Furthermore the definition of information infrastructure appears to put the Federal government in the business of providing the national network -- something that did not appear to be the intent of the original HPCC legislation. The legislation's emphasis on federally created infrastructure makes sense only if one assumes that under a democratic administration, the government will take the responsibility for creating a government funded data network equivalent of the PSTN.

The *COOK Report* could support this as a national objective. The only problem is that the legislation doesn't even begin to address the question of how it would occur. It does state that all these important tools will be developed by agencies whose action will be coordinated by OSTP. As the spring 1991 Department of Energy critique of the intent to let the NSF and FNC run and implement the NREN has shown, this is a plan not likely to meet with success as these agencies follow their own agendas to chase scarce federal dollars in a time of shrinking budgets. OSTP has no teeth to coordinate anything. What is more, under Republican administrations it has not been especially friendly to the idea of government building "infrastructure" that should be left to the whims of the marketplace.

In the area of K-12 education, medical applications and digital libraries, the National Science Foundation ends up with the lion's share of the funds and the testbed networks. These efforts will put in place a de facto national data network policy. Such a policy will have a major impact on the providers of the PSTN. One had better take great care to see that through "cost sharing" arrangements with the NSF decisions are not made that benefit private networks at the expense of the PSTN.

Bringing the NREN to K-12 - A Mission Impossible for NSF?

The press release for the bill says that new educational technology will be developed for K-12. In saying this it completely misses the point. K-12 has a great deal of technology. It needs to learn how to use and assimilate what it already has.

Such technology applications, the bill states: "shall include but not be limited to the following:

(1) Pilot projects that connect primary and secondary schools to the Internet and the National Research and Education Network to aid in development of the software, hardware, and training material needed to enable students and teachers to use networks to--

(A) communicate with their peers around the country;

(B) communicate with educators and students in colleges and universities;

(C) access databases of electronic information; and

(D) access other computing resources.

(2) Development of computer software, computer systems, and networks for teacher training.

(3) Development of advanced educational software.

(b) COOPERATION.--In carrying out this section, the National Science

Foundation shall work with the computer and communications industry, authors and publishers of educational materials, State education departments, local school districts, and the Department of Education, as appropriate."

The pilot projects described in section (1) are already on going. The legislation overlooks a critical point -- namely how network connectivity can be scaled so that it has any chance of being affordable to the majority of the nation's school districts. Training material again is under development from many current grants. However current activity unfortunately often impacts only the most wealthy schools districts. Further more, because of the NSF's predisposition to impose high [read expensive] technology from the top down in a research-oriented approach that is meaningless to most school districts, the impact of fresh hundreds of millions of federal dollars would be felt by only a fraction of our already best off students.

Cooperation with the other agencies listed would indeed be critical. Indeed it would be the single most important task. Unfortunately the National Science Foundation has never shown any ability to work in a meaningful way with the Department of Education much less state Education Departments and local school districts.

An Experienced K-12 Telecommunicator's Critique

Joan Winsor is a Palo Alto California high school teacher and an avid user of computer networks. Here are her comments on the K-12 portion of the Gore legislation:

What is meant by "connect primary and secondary schools to NSFNET?" If that means providing free access via subsidized 800 lines, and such a connection would permit schools TELNET access to other hosts on the internet this could be of significant value to schools. However, if it does not provide access *without* variable connect costs, or only provides access to a yet-to-be-invented K-12 oriented NSFNET host, this would be of very limited value to K-12.

The K-12 arena already has a number of existing "grass roots" mini-networks many of which have already managed to get themselves connected in some way to each other via the internet. Rather than re-inventing the wheel, such a proposal should support and expand existing efforts and provide opportunities for these already-experienced communities to improve and expand their offerings.

Does the bill specify for what kind of computer hardware it proposes to have the NSF develop software? If only one type is specified (no matter what platform is chosen), a majority of the intended "audience" will be excluded because of the wide variety of equipment currently in use in the schools, and the lack of money in school budgets for purchasing new equipment, and/or replacing existing equipment and associated software with another type. What does the High-Performance Computing Program have that mainstream K-12 (as opposed to a handful of special subject area secondary teachers) wants? Not a whole lot.

After showing them what is available on each, ask a teacher whether they would rather be able to connect to NSFNet or to KIDS-93 (listserv), or KIDSNET (listserv) or FrEdMail, or Fido's K-12Net, or Iris, or National Geographic's KIDNET, or BreadNet, or TERC or one of the many other education oriented, educator created electronic networks, I'll bet that NSFNet would run dead last (unless access to NSFNET includes "free" access to TELNET.)

There is nothing interactive there (for education) that I've been able to discover, and teachers do not just want to search databases or download lesson plans or software; they and their students want to interact on curriculum oriented topics. Connect them to the internet without the threat of the "ticking phone meter," and watch what happens. Projects will develop themselves. K-12 does not need NSF to develop educational software. There are plenty of software providers already in the market.

If they have the development of a single internet interface in mind, they are going to need a lot more money than they are proposing to do that - they'll either have to pay to develop it on all platforms currently in use in the schools, or else pay to upgrade all schools hardware to whatever they decide is the standard, or else they are serving only a tiny portion of the population.

There are some schools that are making use of the high end capabilities, but what percentage do you estimate that they are of the total K-12 population.? Far less than 1%, I'm certain. What is needed at this point in time is *access* for the majority, so that they can become involved - at the *beginner* level. The actual percentage of schools who have even put their toes in the water is still very low. Might even still be in the single digits percentage-wise when looking at the entire population. The "high end" folks will continue to find ways to do their advanced stuff - parents and teachers are putting up their own money to make it possible. Teachers who are not permitted to have phones in their classrooms are taking files home and transmitting from their own phones at night. In order for *all* schools to benefit, ACCESS at the low end has got to happen, and the NREN *could* offer such access if the legislation were properly written. Unfortunately, it's a lot more "sexy" to write up a high end proposal, and have the legislators imagining kindergarteners everywhere romping around on the supercomputer, not realizing the many rungs on the ladder that must be climbed before that can be more than "vapor".

Conclusion

The legislation goes on to describe how tens of millions would be thrown at the development of digital libraries, medical network applications and factory network development. Efforts are already underway in all these fields. Unfortunately the legislation does not appear to recognize this and doesn't show the extent to which the federal effort could unfold in a synergistic way with on going developments.

Some will question whether it makes sense to give money to NSF for applications. The agency's purpose has always been assistance to research and usually to basic research. The NSF structure is keyed to cooperative agreements where there is relatively little explicitly required by the agency. The recipient generally thinks of the money as a grant. There to be spent with considerable discretion and without the specific requirements imposed by a contract. If the government is to build an information infrastructure, the lessons of the current mismanagement of the NSFnet T-3 upgrade show that fiscally responsible use of the funds would necessitate contracts not cooperative agreements that function like grants.

There are enough problems in the implementation of the current HPCC legislation. Senator Gore should focus on understanding and fixing them before, he continues down the earlier path of spending money without feedback on the effectiveness of the expenditures for the stated legislative purposes.

CLIN Revisted - Cont. from p. 1)

sweeping set of related components which, together, comprise what is now more precisely referred to as 'CLIN' - the Community Learning and Information Network. I have also had an opportunity to meet with one group of participants, and assist in brainstorming an interesting Research and Development component of the proposed project.

It is clearer how the organization will shape up, with a national not-for-profit corporation heading things up in Washington, answering to a board of directors which comes from business, government, and education. And a for-profit arm to handle the technological products and services offered by companies willing to make long term commitments to the national education and training program. And, headed up by a Chief Scientist, a sort of advanced research and development arm, involving private and public research and academic institutions and companies, and concentrating on the advanced technical tools and services needed, as well as the new education and training 'models' which have to be supported, if the program is to be cost effective in upgrading the skills of the work-force, government, including community based, military reserves and national guard, and the 110,000 K-12 school district students, operating in a 'shared-resources' mode.

The original information 'delivery' concepts of CLN have changed. Initially, they focused almost entirely on one-way delivery of educational 'information', and major dependence on real-time interactive video-conferencing. Initially they also centered on singular physical learning 'centers', i.e. one or two 50 seat 'classrooms' in each community. This early focus has given way to a much broader concept of multi-modal, multi-media, distributed information and interactive communications systems in which the role of the centers may become more a gateway to services, rather than simply a new version of a classroom. The Washington-centered physical Information Service - which implied a giant computer with all the useful information stored on it - is giving way to the idea of a 'virtual' national center which is more a directory and gateway to information which itself may be stored and kept up in thousands of places like the Smithsonian, Library of Congress, or business professional data bases. So that end using publics can 'find' and retrieve the information through the center, rather than getting it 'from' the center itself.

With that metamorphosis of the concept from a focus on centralized data and volume delivery to community centers, to a focus on a network of centers linked to users anywhere - by modem, by interactive satellite, digital radio, and to useful information, resources, and people who can 'teach' subjects to business and government employees as well as students, the whole idea of CLIN takes on a realistic view, as far as I am concerned, of the way telecommunications technologies can be used in advanced ways to make this a 'Learning Society.'

My participation, along with Dr. George Johnston, Research Scientist from MIT, in a San Diego meeting under the moderatorship of CLIN's Chief Scientist, Dr. Mike Wiskerchen, University of California, San Diego, revealed this emphasis on hard-headed systems and technological research as a basis for the national program.

In my first review of CLIN, I wondered where the specific technological and pedagogical beef was to the program. The beginnings of the answers started coming in San Diego. The key players at the San Diego meeting were, corporate educational (not just public education, but work-force retraining) experts and program managers from Dec, IBM, Sprint, Hughes Network Systems, Apple, high-tech researchers and system engineers from MCC, Livermore Labs, MIT, SAIC, Booze-Allen, educational operatives from University of California, California school districts, and representatives from other small and large operations.

One obvious source of advanced research and development support, and well as a major potential 'customer' of the training and educational services to be delivered to local communities, is the Department of Defense. That is because of the twin realities that the military services, which operate the largest and most advanced 'training' and educational systems in the world, will, in an era of cut-backs will increasingly rely on highly trained and ready reserve and national guard personnel for national security needs. Such personnel live in in local communities, and it is very costly in travel and per-diem expenses to transport and maintain them in centralized places. Better that, using technology, the skills are brought to them where they live, than use the traditional methods of bringing them to large military schools and training centers. And without the huge - billions - in travel costs to DOD.

Additionally, large amounts of impressive scientific and technological resources freed up from trying to outdo the Soviets in weapons systems, can be turned to the new meaning of 'national security' as technological prowess and a highly skilled work force become primary in the global economy. The same Livermore Lab talent that had to match Russian science in making smart bombs, and the Army's technological training systems that taught tank gunners that whipped the Republican Guard in Iraq, can make an enormous, near term contribution to the upgrading of production-line work-force skills in the private sector - that neither our public educational system or business itself is prepared to do on a massive scale in time to make a difference to our national economy. At the same time the reserve and national guard part-time soldiers can get, and keep, current in their specialized military skills which, right down to today's infantry soldier on the battlefield, is as high-tech in its way as Top Gun Navy fighter pilots, combat ships at sea, or Air Force stealth bombers. Simulation devices and scenarios, linked by telecommunications, as much as field training, have become the virtual training grounds of the future.

Many a high-tech corporation which hitherto has made computer, display, communications, control equipment for military use can be expected to offer their abilities to the vast and growing market of 'high-tech education and training.' And be willing to help 'seed' that market by making substantial commitments to funding the initial equipment needs of the centers, as well as send their own employees to such centers to learn the skills needed to work in their evermore advanced-skill businesses. And pay for the training. For, since public schools are not turning out adults with the skills they need, US business has only two choices. Either set up their own internal company schools - at great expense. Or support community centers who can train their workers. Note that small and medium sized businesses don't even have the first option.

So the interest in CLIN by business, and the leadership the US Chamber of Commerce in the program, is fully understandable on many accounts.

So the idea remains that the 110,000

school districts in America, who are so budget-strapped they don't even want to think about acquiring new technological tools to teach tomorrows productive citizens, will benefit by having the CLIN's resources available to them a little or no cost, because both business and government will be willing to pay for the center's operations and do a better job at less cost using the same center's technologies for their personnel, as kids learning the science of global warming, or how to construct a solar-powered engine at a high school.

Controversial ideas? You bet. Future Shock? Yes sir. But while pro-defense and pro-big business ultra-conservatives may argue with anti-defense and anti-business politically correct ultra-liberals, and dominate the headlines in their disputes, pragmatic Americans who just want things that work, and who outnumber both political wings, may just take to the 'Community Learning and Information Network' idea, which addresses practical problems of education, business, defense, and taxes in their own communities, and make it work. No guarantees. But anyone have a better idea?

Will Internet Commercialization Soon Include Audio?

About two weeks ago, at a mid-level network, we watched a demonstration of live audio conferencing over the internet. The local speaker had a microphone taped to the front edge of the monitor of his Sun workstation. A window on his monitor showed the names and network addresses of about ten participants in the conversation. The name and address of the person speaking was shown in inverse video. A vertical bar pulsed up and down in modulation with the voice.

I mentioned this to a friend who suggested that I contact Bob Cowan, President of American Teleconferecing Services, number 157 on the INC 500 list of the fastest growing small corporations. ATS was founded in 1984. In 1987 its revenues reached $1 million, in 1990 $3 million and in 1991 $6 million. ATS offers custom conferencing services for regular clients. These range from sales meetings and product introductions, to press conferences, to emergency services that can patch together the top 50 executives of a client company at any time day or night in under one hour. ATS has won an extensive clientele by developing and delivering a broad range of customized services.

Asked what applicability does TCP/IP and the internet have to the the market he serves, Cowen replied that it had incredible potential. Any company with a private WAN could be a user of completely digitized conferencing services as could any member of the current internet. He sees the multiplexing of voice, video and data as one of the next technical frontiers for ATS to address. Use of internet technologies would permit slow scan video to be added to his already formidable arsenal.

Cowan has just hired Paul Christiani as his VP of Applied Technology. Paul engineered and ran the data voice and video net used in Operation Desert Storm. Internet service providers who might be interested in being introduced to ATS should contact Gordon Cook.

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