The COOK Report on Internet → NREN

IBM's Agenda Appears to Become Clear: Therefore It's Time to Put the NSF on Notice! -- An Editorial

Vol. I, No. 8, November 1992

The Missing Pieces of the NREN Puzzle

Michael Miller's article in the September 16, 1992 *Wall Street Journal*, "IBM Commits More Than $100 Million on Venture to Relay Video, Other Data," appears to contain the missing elements of the NREN puzzle. According to Miller, IBM has made a major commitment to a national broadband data network that will, in its first three years, be used primarily for business purposes such as enabling users of corporate LANs to do video conferencing across a Wide Area Network. After that IBM's agenda is nothing less than an alliance with the cable TV companies to bring video entertainment into the home.

To quote Miller: "in an interview offering details about the still shadowy venture, its chief architect, IBM Vice President Lucie Fjeldstad, said IBM and a small group of partners will launch the company within the next year by offering video conferences and other services for business. . . . ."

"Ms. Fjeldstad said she proposed the information network to IBM's board in January 1991 and this May received a green light to find partners and launch a new company. She said that IBM agreed to spend 'in the low hundreds of millions' on the venture in its first five years. . . . ."

"'I've got my dowry, I've got my dad's permission,' she said. 'All I need to do is get married. . . . .'"

"While Ms. Fjeldsted said that IBM's plans will ultimately involve services via cable to homes, she said she wants to launch the new company with a more conventional set of business services that won't involve cable. She said that these would include offerings to let office workers set up video conferences over their PCs and work together simultaneously on linked PC screens."

"In her vision of IBM's strategy, IBM would lure cable companies and other partners by promising to deliver business customers and offering partners a share in the revenues of these initial office services. The joint venture would be expanded into homes within three years."

"Ms. Fjeldsted said she didn't expect the new company would be profitable for at least four years. She also said IBM is prepared to launch the company into the business market on its own if it can't line up any partners."

"IBM thinks its sweetest bait for potential partners, in addition to its clout and cash, is a set of technologies it refers to by the code names Planet, Orbit and Comet. They do the work of managing and distributing the network's data. . . ."

The article concludes with the assertion that "IBM has approved sending an 80 person team of IBM network experts to the new joint venture." It also mentions tests of the new technology with Rogers Communications, Canada's largest cable - TV provider and BellSouth Corp. (These tests of the Paris/Planet technology have been public knowledge for about a year.)

IBM Staking its Future on Data Networking?

When one considers the IBM investment in Prodigy, and the brand new business outsourcing network company called Advantis and adds to it this new multi-hundred million dollar investment, one has to be asleep not to think that something significant is going on.

The *COOK Report* contends that a massive part of what is going on is found in IBM's growing interest in the NSFnet backbone, which to many is the national infrastructure that could tie much of this together. Allowing IBM to insert its proprietary Planet - Orbit, switch and LAN technology into the NSFnet would involve the Bush White House in sanctioning the installation of non standard switching technology into our national "competitiveness" network. If this happened, our switches would speak PTM over Sonet and those of Europe and Japan ATM over Sonet. Where the boundaries met, there would be network gateways that could become massive bottlenecks.

On the Inside

What makes the *COOK Report* concerned? First, after close examination of the draft NSF solicitation in the light of the NASA critique, it appears that the solicitation was carefully crafted to satisfy the interests of IBM. This editorial will state why on a point-by-point basis and put the National Science Foundation on notice that the policy it is following appears to be quite different from what the Congress intended and what the public interest would demand.

If after these points are debated, the NSF does not substantially alter the shape of its final solicitation, it is to be hoped that all the major telcos from the LECs right up to the IXCs will admit the reality of what will have all the appearances of a sweet heart deal and will go to court to prevent the use of public funds in setting up what would be a huge private data network outside the PSTN and outside of international CCITT standards.

If after these points are debated, the NSF does not substantially alter the shape of its final solicitation, it is to be hoped that all the major telcos from the LECs right up to the IXCs will admit the reality of what will have all the appearances of a sweet heart deal and will go to court to prevent the use of public funds in setting up what would be a huge private data network outside the PSTN and outside of international CCITT standards.

Drawing the Link from Fjeldsted to NSFnet

Now readers are probably asking what there is to link IBM's interest in using Planet and Orbit in Ms. Fjeldsted's network with the NSFnet backbone? First ANS will certainly be submitting a response to the final solicitation. Al Weis has made no secret of his intention to use the Planet switch. Secondly Alan Baratz, who has shepherded the development of Planet for several years, went public with Ellen Mesmer in a *Network World* article on July 27, 1992. Mesmer's story is headlined: "Canadian trial could be key to IBM's NSFnet bid." The first paragraph states: "IBM is preparing to test gigabit speed network technology here that might become part of the new backbone if IBM catches the golden ring on the upcoming National Science Foundation contract." The article calls Planet by its old name of Paris and says that it is a router not a switch but it makes very clear that the trial mentioned are the banking trials with Rogers Communications that were mentioned in the *Wall Street Journal* story.

The article concludes: "while IBM appears to be busy with futuristic networking, the company is also engaged in behind the scenes negotiations to create alliances for bidding on the new NSF backbone net. Baratz said that IBM may decide to bid outside of Advanced Network and Services Inc., the IBM, MCI, Merit Inc. consortium."

So IBM may bid on its own. ANS certainly will bid on its own. And IBM may also bid as a partner with another group or groups. To anyone not familiar with the strategic importance of the network as a testbed for IBM's exploration of ways to save SNA and its mainframe business, IBM's infatuation with the backbone rebid might be hard to comprehend. Is it going too far out on a limb to raise as a hypothesis that IBM's financial viability over the next 5 to 10 years depends on salvaging what it can from its mainframe business while at the same time figuring out how to take advantage of the distributed computing environments that networks permit -- the very environments that are threatening its mainframe profitability?

PROBLEM ONE: The NSF draft solicitation, as presently crafted, makes it possible for IBM to walk right in the "open door" with its proprietary switch, LAN and router.

The final parts of the NASA critique, quoted elsewhere in this issue, show both the fiscal, the technical, and the management irrationality of the plan that the NSF has put forth to let the vBNS provider be responsible both for the substrate (levels 1 & 2) and the routing (levels 3 and 3++). Structuring the solicitation this way requires the splitting of the Routing Authority from the vBNS provider to ensure fairness. This in turn increases the size and complexity of the forces involves and creates a situation where, to get what is wanted with limited money, the NSF must cost share by allowing the vBNS provider to sell surplus bandwidth in a situation where from day one the NSF is insisting on having anywhere from 3 to 10 times the bandwidth that critics like NASA say it can use.

Now if the network were divided at the substrate level using Sonet framing and ATM cells as NASA and MCI suggest, this would mean the elimination of the vertical integration that the NSF seems to favor. But note that this is the very integration that leaves the backbone provider with so much power that a Routing Authority must be split off and a complicated, divided and expensive management structure installed. And note that the NSF over the objections of Robert Aiken, its own NREN planner, refused to require ATM as a part of the network switching technology in the draft solicitation. No doubt this was just a coincidence.

Why not eliminate the vertical integration? The rationality and common sense of eliminating seem overwhelming. It is clearly a part of Congressional intent that the network be built in accordance with international telephony standards. The answer appears to be simple: if the door is to be kept open for IBM, two things must happen. First the vertical integration must be maintained because smooth integration of the Planet proprietary PTM technology will *demand* that the awardee have control of level 2 (the Sonet level) as well as control over the higher levels. The second reason leads us to problem two.

PROBLEM TWO: The NSF solicitation, as currently crafted, permits vertical integration that leads the NSF into a specious decision to permit shared commercial use of the backbone -- an action that has already been shown leads to a tilted commercial playing field.

Consider what is happening: the NSF proposes to use tax payer funds to buy far more bandwidth than it can justify using in an architecture that appears to be needlessly expensive and, as a result, has to be shared by allowing the provider to use an extra bandwidth for what ever purposes it so chooses.

Combine this with IBM's newly announced business video data network using the same technology that IBM wants to use on the backbone, and one should begin to see the "fit". It would appear that IBM stands to pick up a $25 million dollar subsidy for doing what it is already committed to doing -- namely develop a

high speed backbone for a multimedia network. Now given the hundreds of billions that disappeared in the savings and loan scandals an extra $25 doesn't seem like much. Nevertheless consider how this leads to problem three.

PROBLEM THREE: It is very unlikely that IBM will draw the attention of the ordinary politician or network user to the "synergy" (given Laurie Fjeldsted's new national network company) inherent in the permission to reuse surplus bandwidth. If the past is guide to the future, what it and the NSF will do is brag about how well the public has been served by IBM in kind contributions that total depending on whose figures you listen to anywhere from 3 to 8 dollars for every Federal dollar spent. It seems to the *COOK Report* that given the conditions established, neither the NSF nor IBM should attempt to claim cost sharing given the synergy between the backbone plans and what IBM is publicly committed to doing anyway.

The hype over the wonderful benefits of leveraging federal money by cost sharing tends to obscure problem four -- the viability of the network as a competitive tool.

PROBLEM FOUR: The motivation to cost share leads to favoritism that is shown to large companies with deep money pockets like IBM. The taxpayer investment in the network is supposed to increase American competitiveness by priming the pump for the even more rapid development of our world class technology in this area. But router and switch development is *not* an area where IBM is noted for its technology competence.

Brocks Meeks' July 7th story in *Communications Daily* brought home very hard what some might even go so far as to call the rank incompetence displayed by IBM in the area of router development for the current T-3 network. To the contrary, state-of-the-art competence in switch development appears to belong the Japanese. Router development, for the current networks under discussion, is however much more important than switch development. Here technology leadership is indisputably held by several small and quite new American companies like Cisco and Wellfleet -- companies with gross revenues of between $100 million and $400 million per year. These are the companies with demonstrated technical prowess that should stand to benefit from the NREN program but will *not* benefit if the spurious cost sharing philosophy of the NSF is not abandoned. They simply don't have large enough cash flows to "invest" $25 million dollars a year in a national testbed.

Thus it may be argued that the stacking of the draft solicitation deck by the NSF in favor of IBM appears to run counter to the Congressional intent of increasing our technology competitiveness.

PROBLEM FIVE: The current structure -- given the size and single minded pursuit of network expansion on the part of IBM as the largest player -- may be faulted for concentration of too much power at key spots. On the one hand the vertical integration allowed the vBNS provider that opens the door for proprietary technology in a publicly funded enterprise, and on the other hand we are confronted with the possibility that the Routing Authority might be granted a monopoly that would control the costs of connection to the network. One had better be very certain that if IBM ended up as the vBNS provider, it could not exert *any* influence on or control over the Routing Authority which by misuse of its pricing power could help to stifle competition for IBM.

PROBLEM SIX: Under the current structure there are probably too few clear boundaries between the key players and their technology goals. One suspects that they tend to define and the NSF tends to rubber stamp. In this area one has to raise the question of the necessity of the provision of video by commercial network service providers. NASA could find no good reason for its inclusion and lots of technical problems that inclusion would raise. On the other hand, given Ms. Fjeldsted's interest in video for her network, it is tempting to surmise that IBM may want to lay down conditions where the mid-levels can be turned into conduits for delivery of its business network services to its end users. The mid-levels would presumably get some sort of fees for doing this. This raises the question of whose agenda the mid-levels are there to serve.

PROBLEM SEVEN: Because of the power and the reach of the NSFnet backbone, it is more than just a data network. Rather it will be the flagship standard to which commercial offerings will have to accommodate themselves. The creation of NAPs by the NSF is tantamount to the creation of a national commercial network architecture by an arm of the Federal government. It is an arena that the NSF should enter into at its peril. It is an arena where, once entered, the NSF had better not act in a high-handed and arbitrary matter.

The respondents have warned the NSF about the problems its designs are creating. Unfortunately when Bellcore's Stewart Personik spoke out in May warning that the NSF was designing a national data network infrastructure without the guidance of the FCC and telecommunications committees of Congress, the NSF thought he was exaggerating the situation. Now Brian Kahin of the Harvard Kennedy School of Government has added his voice to the chorus of concern: "I have observed over the past year a growing reluctance on the NSF's part to acknowledge the critical role it is playing in shaping a broader infrastructure. . . . the [solicitation] reveals how much is left unaddressed."

Decision Time for the NSF

It is time to throw down the gauntlet to the NSF so that four to six months from now, when it issues its final solicitation, it cannot claim that it didn't understand or wasn't warned. It stands at a fork in the road. So far, perhaps out of ignorance, it has taken the path that will serve certain narrow corporate interests. However, it is now known exactly what these narrow interests are. NASA has painted very strongly the technical and management problems that pandering to these interests invites. This essay paints some of the broader policy implications. Others will begin to ask questions between now and the end of the year.

The NSF must now decide whether to proceed with its current direction which appears to be capitulation to the interests of a very large and powerful computer company or whether it will respect the wishes of the larger network community. The stakes are high and the players should send the NSF unignorable messages. FARnet and the telephony community should both publicly warn the NSF that they will challenge the solicitation in court if the NSF does not remove the biases

that have been identified. Should the NSF fail to act in the larger public interest, Congress may hopefully also begin to understand and, starting with the Boucher subcommittee, should undertake to call, what will have become, a renegade agency to account.

The Problems: A Summary

PROBLEM ONE: The vertical integration of the NSF draft solicitation, makes it possible for IBM to walk right in the "open door" with its proprietary switch, LAN and router.

PROBLEM TWO: The vertical integration contained within the NSF solicitation causes the NSF to make a specious decision to permit shared commercial use of the backbone -- an action that has already been shown to lead to a tilted commercial playing field.

PROBLEM THREE: It is very unlikely that IBM will draw the attention of the ordinary politician or network user to the "synergy" (given Laurie Fjeldsted's new national network company) inherent in the permission to reuse surplus bandwidth.

PROBLEM FOUR: The motivation to cost share leads to favoritism that is shown to large companies with deep money pockets like IBM. The taxpayer investment in the network is supposed to increase American competitiveness by priming the pump for the even more rapid development of our world class technology in this area. But router and switch development is *not* an area where IBM is noted for its technology competence.

PROBLEM FIVE: The current structure -- given the single minded pursuit of network expansion on the part of IBM as the largest player -- may be faulted for concentration of too much power at key spots.

PROBLEM SIX: Under the current structure there are probably too few clear boundaries between the key players and their technology goals. One suspects that they tend to define and the NSF tends to rubber stamp. In this area one has to raise the question of the necessity of the provision of video by commercial network service providers. NASA could find no good reason for its inclusion and lots of technical problems that inclusion would raise.

PROBLEM SEVEN: Because of the power and the reach of the NSFnet backbone, it is more than just a data network. Rather it will be the flagship standard to which commercial offerings will have to accommodate themselves. The creation of NAPs by the NSF is tantamount to the creation of a national commercial network architecture by an arm of the Federal government.

Commercialization and COREN in the Context of the Backbone Rebid

Arlington VA. September 14 - 16th, 1992

The CIX and Growing Irrelevance of AUP

After three days discussion at a Tele-Strategies sponsored conference on the Commercialization of the Internet, the status and general shape of the net's commercialization is becoming clear. According to PSI's Bill Schrader 65% to 70% of the Internet is reachable through the members of the Commercial Internet Exchange (CIX). For the thousands of organizations and millions of users involved, the NSFnet backbone Acceptable Use Policy is irrelevant. JVNCnet is just the most recent network that has applied to join the CIX. As CIX membership continues to grow, the irrelevancy of the both the AUP and the NSFnet backbone will increase.

In the meantime ANS and the CIX appear to be in a race to sign the mid-levels. ANS is pushing the mid-levels to sign its "agreements" which give the mid-levels the right to send commercial traffic across its commercial backbone which physically is the same thing as the NSFnet backbone. ANS states that about half the mid-levels have signed the agreements, but does not appear to be willing to state publicly which and exactly how many mid-levels it is talking about. An ANS client of a mid-level that has signed the agreements was told by Joel Maloff at the meeting that if it signed with ANS as a commercial client, it would get a server with routing tables to the mid-levels and find out at that point which it could get to and which it could not. (Upon further questioning about the apparent incongruity of having to pay a $600 surcharge for commercial use and being able to reach *fewer* potential customers as a result, Joel said that the potential client would be told in advance.)

Since the CIX publicly announces its membership, the extent of connectivity being acquired by members is presumably more clear to potential CIX commercial users. While ANS declined to answer the *COOK Report*'s request for the names of the mid-levels that have signed its agreements, it seems clear that, at this point, a company seeking the broadest commercial access will find it through the CIX. (NEARnet, BARRnet, OARnet, CONCERT, CICnet, and NYSERnet are generally publicly acknowledged to have signed the ANS agreements. Some other mid-levels may have signed.)

To confuse matters even more, a mid-level employee who is very knowledgeable about routing (and whose net-

work has *not* signed the ANS agreements) stated that he finds that his network's commercial customers can now use the T-3 backbone to reach the CIX. Traffic outbound goes from his network to ANS, from ANS to the CIX interconnect and from the CIX interconnect to PSI. The return route appears to be PSI -> ANS -> his mid-level. He says this routing only recently became available and he is quite pleased to find that it works.

And in mid September BARRnet introduced for its members a new "Expanded Routing Service" which permits full commercial use of BARRnet and access to the other commercial networks including ANSnet. To cover the cost of the ANS settlements and other overhead of routing changes, members who desire this service will pay a surcharge equal to 15% of their BARRnet annual usage fee.

Waiting in the Wings: COREN

This picture, already murky, may soon become even more clouded before things shake out. The Coalition of Research Networks (COREN - first identified in the May issue of the *COOK Report*), and a potentially extremely important development, is on the horizon. One source states that COREN (which has been the subject of at least one MCI and SURAnet sponsored non disclosure meeting) may be announced with thirty days while another holds out a time frame in the 3 to 4 month period.

COREN, as reported in our May issue, is under the leadership of MCI and SURAnet. It appears that it will be formed as a for profit corporation with, according to one source, as many as three fourths of the Mid-levels connecting to a backbone provided by MCI. (It is reported that MCI and SURAnet have invited AT&T and Sprint to join in the discussions). Needless to say such a backbone will be AUP free.

While such a backbone would likely start at T-1, it would be surprising not to see it move rapidly to T-3 using Wellfleet routers which have been incorporated by MCI into almost every aspect of their networking operations except for the ANSnet backbone. The first source states that "the feeling is that the NSF just can't solve the problems and that ANS is a threat to the infrastructure the mid-levels have so laboriously built -- so it appears likely that most mid-levels may desert the NSF and the NSF backbone."

One source states that COREN (which has been the subject of at least one MCI and SURAnet sponsored non disclosure meeting) may be announced with thirty days while another holds out a time frame in the 3 to 4 month period. COREN is under the leadership of MCI and SURAnet. It appears that it will be formed as a for profit corporation with, according to one source, as many as three fourths of the Mid-levels connecting to a backbone provided by MCI.

A decision by the NSF not to position Network Access Points (NAPs) at all 16 current NSFnet backbone sites would result in additional estimated "back haul" costs to the mid-levels of at least a million dollars a year. If the NSF were to decide to place the NAPs elsewhere, this would be likely to provide the final push in forming COREN or, if it has already been formed at T-1, for pushing the backbone to a true T-3 where MCI, using Wellfleet routers, would be free to give network users vastly improved throughput compared to the ANS backbone that is encumbered by what *Communications Daily* has called the "crippled" IBM RS6000 routers.

MCI Vice President Jerry Edgerton during his public presentation stated that MCI was working with SURAnet and the "emerging COREN group." When I told him what I knew about COREN's plans, and asked for comment, he said I should query SURAnet. (I have done this and SURAnet declines comment.) However, he also added that MCI is willing to work with mid-levels who want to preserve what they have built and show them how to become fully professional self-sustaining organizations -- a statement that would seem to confirm MCI's leadership role in a potentially beak-away organization that may render the NSFnet/ANSnet backbone meaningless if the break-away mid-levels disconnect from it.

If the NSF agreed to continue to underwrite the full cost of mid-level connection, they probably would not break away. The function of COREN at that point might be tantamount to that of the commercial network interconnects envisioned by the NSF plan for Network Access Providers. For the NSF, presumably the only problem would be that they would no longer be in control and that their plan for NAPs could be rendered superfluous before being put into effect. Edgerton said that the NSF had some interesting challenges to think about over the next 3 to 4 months and that the final development of COREN would be clarified within that time frame. Another source thinks that the fall out will come sooner -- probably within a month. Clearly MCI is ready to compete with ANS. Reading its separate response to the NSF backbone solicitation draft makes this very clear.

Certainly commercialization of the Internet is here to stay. Maloff for example pointed out that he estimates the current world wide market for TCP/IP hardware, and software and associate network services at $4 billion a year. Others added that 90% of the Fortune 1000 are expected to join an AUP free net (either the CIX or COREN) within three years. This does not mean that they will dismantle their private networks which they will continue to use for their proprietary most sensitive data. What they are likely to use the commercial Internet for is electronic mail -- including such areas as communications with field service crews. Proprietary data will come much slower. In the meantime people at companies like DEC are thinking of beginning to develop software including security tools that will help to facilitate their entry.

Certainly commercialization of the Internet is here to stay. Maloff for example pointed out that he estimates the current world wide market for TCP/IP hardware, and software and associate network services at $4 billion a year. Others added that 90% of the Fortune 1000 are expected to join an AUP free net (either the CIX or COREN) within three years. This does not mean that they will dismantle their private networks which they will continue to use for their proprietary most sensitive data.

Responses of Key Players to Flawed Backbone Solicitation Place NSF on Notice

The NSF Stands at a Fork in the Road

The National Science Foundation would appear to be on one of two tracks. It may have a master plan for continuing the network in the hands of ANS or an ANS-like entity and may be willing to take the potential legal and political flack from such a course of action. On the other hand, the NSF may be genuinely trying to learn how to redo the network in the best public interest. The *COOK Report* hopes that the second possibility is the reality. If so, given the withering and often contradictory criticism directed at the Foundation in the 232 pages of comments from 44 entities (and increasingly in the computer press) we predict that the NSF is likely to be another six months away from issuing a final backbone solicitation. However, when the final solicitation eventually does come out, the NSF will very likely be making its choice between the paths very clear.

Several respondents suggested that the NSF actually come out with another draft for comments before going to the final solicitation. But at the Telestrategies meeting on September 16, George Strawn the NSF solicitation manager said the NSF would *not* do this. There is reason for concern. The Foundation's ability to cope with the task that it faces - one of creating a fiscally *and* technologically sound backbone infrastructure - given its propensity to embrace the NREN "hype" of buying unneeded artificially high speeds from day one, may mean that it develops a monster which will needlessly gobble up all the extra NREN moneys in new subsidies for network connectivity in a more expensive monopolized environment.

Given the scope of the task, the money at stake, and the apparent direction to be taken by the NSF, sound policy analysis would raise the question of whether the NSF has anything like the resources needed for discharging its intended program in a fiscally responsible way. The criticisms levied make it plain that the way it is approaching its task is likely to create another contentious situation that, with no performance criteria, is likely to result in money being thrown at problems rather than being used in fiscally responsible ways.

The ANS Response

The ANS critique was extremely detailed and self-serving. The general view is nicely summarized in the statement that the response is offered so that the "NSF goals for advancing research, education, and US competitiveness can be realized." Tellingly absent is the goal of ubiquity of connection to the network.

The response underlines ANS' requirement for being able to sell backbone bandwidth for commercial use in view of the fact that the price the NSF is willing to pay will enable "bidders to recover only a fraction of the real costs required to provide pre competitive vBNS service at multiple network access points (NAPS)." ANS states that it "will be necessary for the NSF to promote the vBNS as a shared service to attract qualified providers to bid on the NSF solicitation." When the NSF from day one decides to buy perhaps ten times the bandwidth it can reasonably be expected to use, it does increase its cost for sure. As NASA points out there are other more responsible ways to proceed.

ANS asks the NSF to tighten rather than loosen AUP by requiring "an end user service subscriber to be identified by a unique network identifier" that would allow the network service provider to impose an "AUP administration" on that user. The problem with this is that commercial subscribers (many of whom are small information age businesses) would find themselves increasing pressured to sign up as commercial customers at ANS' higher commercial rates. ANS' desires appear to be to move the burden of AUP compliance from the end user to the network service provider that provides that end user with network connectivity.

ANS does say that "a standard administrative mechanism must be developed to ensure that NSF funds are not used to directly or indirectly subsidize commercial vBNS services." A desirable goal that appears to be missing from the present contract. However, the *COOK Report* points out that the current solicitation that starts out by buying, from day one, more bandwidth than is needed with an absence of any measurable performance criteria for the supplier represents, on these grounds alone, a significant subsidy of commercial services.

On the one hand ANS says "We still support the draft solicitation's position that the routing authority and vBNS provider should be non overlapping organizations or consortia." While on the other it its next sentence states: "Potential NAP manager awardees should include competitors for the routing authority and the vBNS service awards. NSF should encourage all bidders to submit creative proposals to provide NAP manager services." Quite a juxtaposition between the two sentences! It sounds like ANS is saying that the vBNS awardee could also become the NAP manager which is the entity that sets prices for attachment to the backbone. This would create a gross conflict of interest that one would hope that the NSF would not allow itself to be pulled into.

"Potential NAP manager awardees should include competitors for the routing authority and the vBNS service awards. NSF should encourage all bidders to submit creative proposals to provide NAP manager services." . . .It sounds like ANS is saying that the vBNS awardee could also become the NAP manager which is the entity that sets prices for attachment to the backbone. This would create a gross conflict of interest that one would hope that the NSF would not allow itself to be pulled into.

Nevertheless ANS goes on to try bolster its argument by saying that "finding a single NAP manager that is viewed as neutral and unbiased by all network service providers will be difficult. For these reasons, we strongly recommend that the solicitation criteria for the NAPs emphasize the facility, cost, and operational requirements rather than the need for organizational neutrality since this will be a difficult criterion to satisfy." Since ANS has never been a fan of level playing fields, its views are not surprising. We add that a requirement for standard pricing and access conditions on the part of the NAP manager is not very helpful since under the conditions ANS is proposing all this would mean is that the vBNS awardee, if it held both contracts, could set them in a uniformly burdensome manner.

In what some might interpret as a move to weaken the mid-levels by driving up their costs, ANS adds that "we believe the best choice for NAP locations may be a carrier's point of presence. These facilities provide for cost effective co-location of networking equipment and infrastructure."

In what some might interpret as a move to weaken the mid-levels by driving up their costs, ANS adds that "we believe the best choice for NAP locations may be a carrier's point of presence. These facilities provide for cost effective co-location of networking equipment and infrastructure." In a further move that seems designed to ensure that the backbone provider could control the attachment costs of its commercial competitors, ANS goes on to say: "A quick inquiry among the four largest US inter-exchange carriers would demonstrate that there are already several common long distance inter-exchange carrier meet points that could support a NAP. It might be cost effective to choose NAPs that are co-located with same inter-exchange carrier that supports the vBNS service since there will be no vBNS local access costs."

Finally it appears that ANS does not want the routing authority to have much authority. "Where there are disputes involving the peering relationships between the different NSPs, the routing authority should refer the disputes back to the participating NSPs for resolution. The routing authority should not be responsible for the enforcement of routing policies or the NSF AUP." One wonders what the routing authority WILL be responsible for?

MCI Prepares to Go its Own Way

Meanwhile MCI submitted its own set of comments on the solicitation. The fact that it chose to comment independently of ANS, and that its content is dramatically different from the NSF design intention that favors ANS, leads the *COOK Report* to believe that it is prepared to submit either alone or with Wellfleet, its own proposal to deliver the vBNS. What MCI said makes good policy sense. It is essentially in conformity with the way that NASA and DOE have done their backbones with Sprint.

Meanwhile MCI submitted its own set of comments on the solicitation. The fact that it chose to comment independently of ANS, and that its content is dramatically different from the NSF design intention that favors ANS, leads the *COOK Report* to believe that it is prepared to submit either alone or with Wellfleet, its own proposal to deliver the vBNS. What MCI said makes good policy sense. It is essentially in conformity with the way that NASA and DOE have done their backbones with Sprint.

"MCI proposes a modified structure . . . [that] would bring about dramatically different relationships between the providers. MCI submits that the responsibility for layer 1 and layer 2 service should be separated from the responsibility for layer 3 and 3++ service. [Under the NSF draft] the vBNS provider is responsible for TCP/IP routing in the vBNS. MCI submits that the responsibility for the IP/CLNP routers connecting to the backbone be assigned to the NAP manager/routing authority. . . . The routing authority not only manages the route servers at the NAPs but also the vBNS router at each NAP. This structure limits the vBNS bidders to providing layer 1 and layer 2 services only. Typically this is what the IXCs and LECs have announced to be their next generation of services (SMDS/ATM). Thus the NSF would leverage their strengths and interoperability at this level."

In a section called: "benefits" MCI states that such a "strategy aligns the vBNS with the direction of the communications industry. Doing so accomplishes several objectives: It positions the NSF Network for a most advantageous growth and cost sharing opportunity within the industry. It increases the competitive nature of the responses to the solicitation since the costs of providing the basic backbone services are no longer obscured by costs associated with packet routing and switching functions at layer 3 and above. It creates an environment that is indisputably a 'level playing ground' at layers 1 and 2. This environment is most likely to bring out the lowest cost alternative."

Finally MCI points out that if the NAP Manager/RA controls not only "the route servers at the NAPs but also the routers that connect the vBNS to the NAPs, that is the IP/CLNP routing on the vBNS. This will give additional focus and control to the routing authority function. It also eliminates potential conflicts between the NAP Manager Routing Authority and the vBNS provider." The COOK Report believes these are sound judgments where a participant in the current backbone mess is actually trying to learn something from past mistakes rather than repeat them.

Bellcore Chimes In

Bellcore submitted a succinct set of comments on behalf of most of the RBOCs, five of whom also commented separately. Rather than divide functions it recommends dual networks as did the first issue of the *COOK Report* in April. It suggests that the NSF:

"Adopt a structure composed of Production and Experimental Networks as a means of addressing the issues raised above. The Experimental Networks would constitute a government supported testbed for network functionalities that would support the Grand Challenges. The Production networks would evolve from the current three tier hierarchy based on the NSFnet and would be managed and operated by multiple providers from the private sector."

"Limit access to the Experimental Networks to only those researchers directly involved in the technological experiments performed and research applications developed on those networks. Procedures to ensure this objective should be made a requirement for all responses to NSF solicitations."

"Utilize Production Networks to carry Research and Education traffic that currently flows on the NSFnet (i.e. traffic that meets the current NSFnet AUP), but which is not related to the testbed mission of the Experimental Networks."

"Directly subsidize end users in the Research and Education community for their access to and use of the Production Networks."

Bellcore's comments represent a very reasonable approach designed to assure a level playing field and avoid the government's picking a winner by selling a market place advantage to the highest bidder. However NASA went Bellcore one better with the depth and specificity of its critique.

NASA Unleashes a Withering Barrage

The very long and detailed NASA critique submitted over Milo Medin's signature begins by noting that the NSF has set a lot of requirements without explaining *why* it wants them. It recommends "that discussion be added to each area where requirements are being outlined, which gives the background and need associated with each requirement, along with an overall priority to the Government for each requirement. At a minimum, the requirements should be grouped into a mandatory and desirable category, as was done for the NSFnet NIC solicitation."

NASA has great difficulty in finding the speed of 155 Mbs for the vBNS to be justified. "It appears that the vBNS is supposed to be a high performance state of the art backbone network, but that many organizations may not use it, thus causing the 155 Mbps requirement to be called into question. Also, very few of the present mid-level and regional networks have links in their backbones that exceed T1 rates. It is questionable that even if the backbone provides 155 Mbps service for aggregation, the served networks will effectively be able to use it."

"In the FY 94 time frame, many of the major inter-exchange carriers or IXC's, will not be in a position to support Sonet capacities of this magnitude on a cross country basis. Thus, by requiring 155 Mbps service if it is not absolutely imperative in this time frame, NSF may exclude a large number of IXC's, which will unnecessarily limit the Government's options for award."

"In the FY 94 time frame, many of the major inter-exchange carriers or IXC's, will not be in a position to support Sonet capacities of this magnitude on a cross country basis. Thus, by requiring 155 Mbps service if it is not absolutely imperative in this time frame, NSF may exclude a large number of IXC's, which will unnecessarily limit the Government's options for award."

"Thus, while we believe it is clear that at least a true DS-3 based network is required, making the jump to 155 Mbps at least initially does not seem justified by the overall structure of the architecture of the network and it's intended uses. NASA recommends that the solicitation better identify the requirements for greater than DS-3 connectivity, and state the mandatory requirement as being a true 45 Mbps or above, and an eventual (possibly the FY 95-96 time frame) highly desirable option for 155 Mbps and above service."

No Performance Requirements

Given the past five year history of the current backbone performance problems NASA understandably wonders why there are no performance metrics in the draft solicitation. "Given the history of the existing NSFNET backbone awardee in terms of delivering actual high performance connectivity, the absence of specific performance and reliability targets and objectives seems especially problematic. This problem is also exacerbated by the Foundation's willingness to allow the vBNS provider to use the same infrastructure to sell to other customers on a non-interference basis. In order to tell whether of not the Foundation is getting what was agreed upon, performance metrics must be set so that base performance requirements are understood, and that adverse impacts by other vBNS customers may be identified and corrected."

"Given the history of the existing NSFNET backbone awardee in terms of delivering actual high performance connectivity, the absence of specific performance and reliability targets and objectives seems especially problematic. This problem is also exacerbated by the Foundation's willingness to allow the vBNS provider to use the same infrastructure to sell to other customers on a non-interference basis. In order to tell whether of not the Foundation is getting what was agreed upon, performance metrics must be set so that base performance requirements are understood . . .

"Even when the present NSFnet backbone awardee first came on-line with a DS-1 based network, it was incapable of delivering true DS-1 performance. [DS-1 means T-1] For some time, the DS-1 trunk lines were multiplexed into 448 Kbps circuits, which were as fast as the existing router based serial line hardware could support, . . . "

"When the upgrade to DS-3 based trunking occurred, the delivered service could not exceed 12Mbps in the best case on a DS-3 trunk interface, capable of 90 Mbps (45 Mbps full duplex), and the FDDI [100 Mbps] interfaces which were offered were incapable of dealing with offered load[s] greater than ethernet [speeds]. Many of these problems have yet to be remedied, even though the intended end date of the award is imminent."

"The nature of the [performance] criteria does not have to be elaborate. One possible metric is something like the system must be able to deliver at least 90% of the committed capacity 75% of the time, and that the network must be able to deliver 20% of the capacity 99.5% of the time. A Mean Time Before Failure (MTBF) and Mean Time to Repair (MTTR) minimum specification should also be required. These specifications should *not* be developed by the vendor. "

"The Government should set these criteria based it on it's programmatic objectives; that is, based on what it needs to service the research and academic communities, and the vendor could be allowed to propose implemetation of these measurements and engineering of the system to meet the specifications and what remuneration would occur when agreed upon goals were not met."

"This would also help protect the Government from criticism about being taken advantage of. When clear specifications are outlined, all parties understand what is being agreed to. Additionally, if this vBNS service is ever hope to be made available as a commercial service offering, some quality of service specifications must be developed and implemented if it is to be successful. The vendor might as well get used to this from the start, and any vendor who really intends this being a commercial service will be very receptive to such specifications in any case, since it makes clear to them what their performance targets for a satisfied customer are."

"Because of these past problems and the needs of network user communities, NASA strongly recommends language be added to the specification that stresses operational robustness as being critically important, and that specific performance and reliability specifications be developed, commensurate with the actual service requirements of the user community, and added to the solicitation."

Problems with Vagueness of NAP Numbers and Locations

"The solicitation does not outline proposed numbers or locations for NAP's, but states that this information will be present in the final solicitation. It is rather unfortunate that a proposed NAP configuration was not outlined in this draft, since it would have probably given the Foundation more focused feedback from the network and user communities." The NASA response goes on to discuss the development of the current network architecture and by taking BARRnet as a specific example shows the problems that would be created if the NSF should decide to site the NAPs a locations other than the current backbone nodes. The way that NASA puts it is that "more of the backbone costs would be borne by the end sites for less optimal service." Another source was less subtle stating an estimate of a bill to a key group of mid-levels of at least a million dollars a year more than they pay for their current backbone connections.

The NASA response continues: "Also, the issue of redundant connectivity must be addressed by the mid-levels, and many of the busiest ones would require redundancy at the DS-3 rate as well. The likely outcome of all this is that the mid-levels will attempt to link themselves together by "back-door" links, and certainly attempt to cooperate on mutual backup routing. [Editor -- COREN comes to mind.] Since these activities would not be occurring at the NAP's, the RA would be powerless to help, and the track record of many of these organizations installing back-doors and providing backups to each other is far from good. During the time the old DS-1 based NSFnet backbone was having so many problems, the mid-levels attempted to do backup through each other and installed a number of very troublesome back-door links, with the result of routing chaos and even less reliability for the overall system. It was a very bad time in the history of the Internet. A repeat of this experience must be avoided."

"In summary, the locations of the NAP's are critically important to the design of much of the network infrastructure. If the NAP's do not provide service to the existing set of mid-levels at existing backbone sites, many mid-levels will be faced with enormous costs to adjust to the new environment. These costs may not be bearable by many of the mid-levels, which are already under enormous fiscal pressures. This will result in cost cutting measures, such as cutting back on operations and management staffs, redundancy, user services and other services, as well as large increases in subscriber prices. . .

"In summary, the locations of the NAP's are critically important to the design of much of the network infrastructure. If the NAP's do not provide service to the existing set of mid-levels at existing backbone sites, many mid-levels will be faced with enormous costs to adjust to the new environment. These costs may not be bearable by many of the mid-levels, which are already under enormous fiscal pressures. This will result in cost cutting measures, such as cutting back on operations and management staffs, redundancy, user services and other services, as well as large increases in subscriber prices, which would occur at a point in time where Universities are already under severe budget pressure because of state and local cutbacks. Many of these organizations might not survive, and while many institutions will be able to switch to more expensive providers, some will have to settle for lower performance or less reliable access, hurting their ability to meet the networking needs of the scientific community, and agencies like NASA and DOE."

"From a more NASA programmatic view, it should be pointed out that we not only depend on the NSFnet backbone, but also the mid-levels for reaching our Principal Investigator's (PI's) at the various academic and research facilities which the NSFnet system interconnects. Also, as we move to more and more collaborative environment with our scientists, there will be much more direct communications between the PI's. This direct PI-PI interaction occurs primarily over the infrastructure and *not* over NASA specific networks. Thus if the NSFnet backbone solicitation fails to produce an operationally reliable facility, NASA mission specific networks will be very hard pressed to support this PI-PI interaction".

"Furthermore, while NASA (and presumably DOE as well) is capable of adequately bypassing the NSFnet backbone under existing budget projections, we are not capable of bypassing significant amounts of the regional and mid-level infrastructure, as a large amount of DS-0 and DS-1access lines would be very expensive. Thus we are very concerned about the potential effects of a new network architecture on the mid-level communities, as they are key to many of our science programs."

"NASA therefore strongly recommends that NSF site the NAP's at existing NSFnet DS-3 backbone facilities, which support the existing mid-level networks that interconnect the vast majority of academic and research sites, and if cost would prohib-

it connecting all of them at the 155 Mbps rate, then preference should be given to interconnecting all at DS-3 and fewer, if any, at 155 Mbps. We would, at a minimum, urge that a detailed analysis of what the likely impacts on the existing infrastructure would occur by any new set of interconnect points be carried out and discussed in a public forum with the affected parties. If a major change is going to happen, it ought to be planned and well understood, rather than stumbled into."

Network Architecture and Policy Issues

"The new solicitation has a different architecture than what the existing awardee, and other networks have implemented. Architecturally, the separation of the routing configuration and interchange from the packet forwarding function should be straightforward with protocols such as BGP and even with EGP. However, in practice, implementations might run into problems that were unforeseen. Since this is a crucial component of the architecture, it should be prototyped and tested with common router hardware and software for robustness before the solicitation is officially released. Again while things should work straight forwardly, since this is a central part of the design, it really should be prototyped first. In fact it should be relatively easy to set this up in a laboratory environment and then on a testbed with willing subjects first, demonstrating proof of concept."

"NASA thus recommends that the RA/ vBNS/NAP client configuration be prototyped and tested before it is released as a central part of a solicitation such as this one. All major functions should be prototyped and tested this way."

"The Routing Authority is to run a route server at each NAP. Since the route server would be a single point of failure for all packet level interconnectivity at the NAPs, it is likely that redundant route servers would be required. Additionally the route servers would have to be operated and managed on a 24x7 basis because even off hour failure would have serious ramifications on network connectivity. Computers would be required to act as these route servers and repositories of the network configuration information, and would be required to exchange routing information with subscriber networks. However, no mention of the need for operation or maintenance of route server hardware or even the computers and their physical NAP attachments is made in the solicitation. Given that a 24x7 management center requires significant staff and management support, this should be detailed in the specification as a requirement for proper operation. "

"Then there is the issue of proper operations and management of the NAP itself. Since proposers are allowed to bid any type of network for an interconnection, the issue of trying to define the interface between the vBNS and the NAP in the solicitation is impossible. The only guidance provided is that the NAP must be able to operate at speeds of 100 Mbps or greater. Since this could be interpreted as an aggregate number and not a backbone speed, this allows people to bid multiport ethernet bridges, an FDDI ring, a large Northern Telecom PBX, etc. This is unacceptably loose for a component that must be common between two separate entities."

"We suggest that something fairly generic be required, such as an FDDI concentrator or ring, and allow the vendor to optionally propose alternate technologies that the Government would choose to allow to be used. This could be treated as a special use of the general procedures for upgrade and support of the NAP's. Also, as with the route servers mentioned before, only the words 'establish and maintain' are used to describe what the awardee is to do with the NAP. This again leaves out any explicit operational responsibility and performance metrics for the NAPs which are crucial to maintaining the connectivity of the whole Internet!"

"Also, the specification basically gives an exclusive franchise for interconnectivity to the NSFnet backbone to the RA, but does not limit the fee schedules to provide only for cost recovery for operations and management of the NAP. That is, there is nothing to prevent the RA and possible subcontractors from charging extortionate and uneven pricing for subscriber attachments according to the specification. Equitable treatment of various network service providers only extends so far as route administration, and *not* into other areas such as attachment and fee collection. We certainly hope this was an error in the specification and not intended to allow abuse of the Routing Authority's role as the "gateway" to the NSFnet backbone. Language must be added to the specification to limit fees to reasonable cost recovery at the minimum, and to provide fair treatment of subscriber networks, with the possible exception of the subsidization of research and academic networks who merely look at the route server as a means of interacting with the vBNS. The Government must remain in control of its right to access its own facilities. Tight policy control must be exercised over the RA to ensure that proper access to the Government provided services such as the vBNS are maintained."

"Also, the specification basically gives an exclusive franchise for interconnectivity to the NSFnet backbone to the RA, but does not limit the fee schedules to provide only for cost recovery for operations and management of the NAP. That is, there is nothing to prevent the RA and possible subcontractors from charging extortionate and uneven pricing for subscriber attachments according to the specification. Equitable treatment of various network service providers only extends so far as route administration, and *not* into other areas such as attachment and fee collection. We certainly hope this was an error in the specification and not intended to allow abuse of the Routing Authority's role as the "gateway" to the NSFnet backbone.

"NASA strongly recommends that the section that deals with the RA and NAP architecture be rewritten to explicitly deal with the critical operational issues and performance requirements associated with management of the NAPs and route servers. It would be sheer negligence to not call out this critical issue in the specification. Also, the issues of limiting NAP attachment fees to that required for cost recovery must be addressed, and no perception of an exclusive franchise for determining access to the Government provided vBNS be advanced. The Government must remain in con-

trol of it's own assets, in this case, vBNS transport services." Strong language, but given the NSF's willingness to tolerate a tilted playing field with ANS, it seems to be not a bit too strongly phrased. There are strong elements of déjà vu in the NSF draft solicitation. One wonders who's writing the script?

The Government must remain in control of it's own assets, in this case, vBNS transport services." Strong language, but given the NSF's willingness to tolerate a tilted playing field with ANS, it seems to be not a bit too strongly phrased. There are strong elements of déjà vu in the NSF draft solicitation. One wonders who's writing the script?

NAP Subscriber Network Attachment Issues

"The specification calls out a number of requirements for NAP subscriber networks. Unfortunately, some of the reasoning behind these requirements is vague. For example, there is a requirement that NAP subscriber networks must be able to carry CLNP. While this is appropriate for the vBNS, it is no appropriate for subscriber access. There is no reason why an organization who chooses not to support CLNP should be forced to do so. Encouraging CLNP transport is one thing, but requiring it will force all manner of legalistic interpretation. What does "able to switch both CLNP and IP packets" mean anyway? That all of the subscriber's network be able to switch CLNP, including things like SLIP links? That some of it must? Maybe only one link in the network must be able to carry CLNP? How about being able to encapsulate CLNP over IP for that matter? In that case, does anything that carries IP qualify as being able to carry CLNP? Either is means all the network must be able to carry CLNP, in which case the requirement is onerous , or it means that anything will be able to qualify as being able to carry CLNP, in which case why does it belong in the specification at all? What about operating in coordinated NSAP space? What about supporting standard CLNP routing protocols? It goes on and on."

"Much the same set of arguments apply to the operation at speeds of 1.5 Mbps or higher. Does this mean all of the network must operate at DS-1? Part of it? Only the connection to the NAP? At the minimum, this language must be tightened up in order to mean anything. Perhaps what is meant is that the subscriber network's backbone has to operate at speeds of DS-1 or higher? In that case, it ought to be stated as such, though how one characterizes a network's backbone is not inherently obvious."

"An even worse example of a spurious requirement applies to the video reference. What does this mean? Does it mean that the network actually must have video codecs attached to it? That someone can run video over IP on it? If that's the case, since the subscriber network supports IP, does that alone qualify for it being able to support video?

"An even worse example of a spurious requirement applies to the video reference. What does this mean? Does it mean that the network actually must have video codecs attached to it? That someone can run video over IP on it? If that's the case, since the subscriber network supports IP, does that alone qualify for it being able to support video? What about the ability of the network to support multiplexed portions of bandwidth carrying video? Does this mean all the CSU/DSU's in the network are able to support this? Some? What about signaling? And how is this video capability supposed to interface to the video capability in the vBNS? If the vBNS carries video in packet mode, and the subscriber in multiplexed mode, how are they supposed to interoperate? Is that even a requirement? This whole requirement should be dropped."

"Perhaps a much more controversial issue is the requirement to subscribe to the routing policies of the NAP Manager Routing Authority. What does this mean? If it means having a routing table that is the same as the one in the route server, then none of the peer agency networks could be NAP subscribers. For example, NASA has a private T-1 link to the University of Arizona to accommodate bulk data transmissions to a large number of PI's there from a variety of NASA astronomical data repositories. The NSFnet backbone has a path to the U of Arizona via Westnet. NASA does not use that path, but rather the private t-1 path. Does this mean that we would be violating one of the NAP attachment criteria? What does this criteria really mean?"

To sum up: "Organizations will be *very* fussy about requirements for attaching to NAPs. Unless these requirements are reasonable and clearly stated, there will be considerable conflict, most of which will probably be needless. The Foundation will almost certainly give in at some point in time since many of the requirements are specious and do not hold up to close examination. NSF should recognize this and weaken this language to recommendations. This is an area where NSF should be leading, not directing. Organizations have been motivated by incentives in the past and that approach should be continued with this solicitation. NSF will not succeed in playing the global routing policeman by coercive means."

Interagency Issues

"The solicitation makes no mention of how the existing NSFNET backbone interconnections to the FIX points would be structured in the new award. The existing FIX interconnections are critical to the proper operation and engineering of the Internet as a whole, and to the Federal research networks in particular. Concerns about operational stability would certainly demand this issue being explicitly addressed in a manner consistent with the overall objectives and evolution of the other Federal network backbones."

"Our position, as that of DOE, is that FIXes and NAP's are two different types of entities, with different goals and objectives. For example, open access is a clear driver in the NAP case, whereas it is explicitly prohibited at the FIX points. This is true for many reasons. For example, the existing FIX points carry multiple protocols, including DECNET Phase IV, which has no firewalling or filtering capabilities. Allowing open access to the network that is the DECNET IV glue for the DECNET Internet (still critically depended on by many physicists and space scientists) would allow uncontrollable DECNET routing problems to occur if a NAP subscriber deliber-

ately or inadvertently switched on DECNET IV support. Similar issues may result with CLNP routing given the requirements of DECNET V/IV compatibility. Also, the connection of DoD and DARPA related networks which may have privacy and access concerns are also in conflict with NAP open access goals. A NAP subscriber network could go into a bridging mode (without the RA or NAP Manager even noticing and be able to transfer DoD to NASA interactions back to a remote location for the purpose of recording passwords being sent in the clear between federal agency facilities. The existing FIXes, especially FIX-West, which we manage, are relatively secure in these aspects because of limited access to the medium. . . . ."

"Because of the existing structure of each FIX being a single LAN in abox (e.g. multiport ethernet transceiver, FDDI concentrator, etc....), operational management, statistics collection, and support issues are greatly simplified. Since there is no need to support a large group of subscribers, scaling issues are also considerably simplified over the NAP case, thereby resulting in fewer hardware upgrades and system transitions that affect users. Whereas the NAP's must be geographically distributed over a wide area, only two FIXes are really needed, as any more begins to add complications to routing designs, already complicated enough. And because 'fair access' to all providers is not a requirement, locations and evolutions of FIX architecture are simplified compared to the NAP case, which must accommodate a larger dynamic range of network subscribers, without the context of an AUP. . . . ."

"In summary, NASA recommends that NSF explicitly address FIX interactions as separate from NAP interactions in the solicitation, and provide high performance NSFnet backbone access without the requirement for route servers to the existing FIXes initially, and address an orderly evolution of FIX architecture with the other agencies that they interoperate with at the FIXes."

Routing Complexities

"The solicitation outlines a set of topological interconnections that will be very complex to manage and operate in the best case, and will simply breakdown into chaos in the worst case. The architecture is in essence trying to allow a large variety of interconnections between the various mid-levels, commercial network providers, Federal agencies, and of course the vBNS itself. It almost appears schizophrenic in that it wants to allow non AUP restricted transport of data between these organizations connected at the NAP's, but yet also wants to be the general purpose network provider as well. It reads almost as through NSF wants to get out of the network transport business but still wants to build the best and fastest network around. These are not mutually compatible goals and the solicitation suffers because both are attempting to be pursued."

"The solicitation outlines a set of topological interconnections that will be very complex to manage and operate in the best case, and will simply breakdown into chaos in the worst case. The architecture is in essence trying to allow a large variety of interconnections between the various mid-levels, commercial network providers, Federal agencies, and of course the vBNS itself. It almost appears schizophrenic in that it wants to allow non AUP restricted transport of data between these organizations connected at the NAP's, but yet also wants to be the general purpose network provider as well. It reads almost as through NSF wants to get out of the network transport business but still wants to build the best and fastest network around. These are not mutually compatible goals and the solicitation suffers because both are attempting to be pursued."

NASA explains in detail AUP problems in routing between the CIX and BARRnet and CERFnet. It goes on to say: "given the state of routing and forwarding technology that exists today, there are relatively few options for dealing with such problems. One is topological control where interconnectivity is limited in such a way as to preserve policy controls. An example would be not allowing any of the NSFnet mid-levels which primarily services the research and academic community to interconnect to the CIX. Another approach would be to do away with the NSFNET backbone completely, and let the mid-levels fend for themselves. Clearly, none of these is particularly appealing, but without research and deployment of new technologies, this problem will not go away. Encouraging large numbers of networks to interconnect at the NAP's will only make this problem worse."

"The solicitation implies that many of these issues will be fixed by the use of the route servers. In fact, this will not fix the problems, but will result in conflicts as incompatible policies are tried to be enforced by the various subscriber networks. Since the vBNS is allowed to carry other traffic than just NSFnet's, it too will have it's routing complicated by these problems. The solicitation is unclear about what policies the vBNS is supposed to implement, yet it's ability to carry any traffic between the NAP's depends on what exactly these policies are. By taking a hands-off approach, the Foundation is inviting more of the same trouble it had with the decision to allow ANS to "sell" access to the NSFnet DS-3 based system, or at least the perception it was doing so."

"NASA recommends that NSF reevaluate it's decision to make the NAP's AUP free. We believe that this approach was based on a mistaken impression that this would solve more problems than it creates, and instead refocus itself on providing robust, high performance network connectivity to the research and academic community. It should also take steps to accelerate work in thefield of source specific routing and policy based routing in general, and try to motivate organizations to deploy said technology. We believe the best way of dealing with the AUP issue is allow connection of commercial nets to the NSFNET backbone system, but to exchange only AUP compliant traffic, and to not encourage mid-level organizations to build dual use systems which are effectively impossible to route properly with current technology." NASA seems to be taking AUP requirements seriously, an area where the NSF says it takes requirements seriously and then appears to approve unenforceable situations."

Program Management - The RA-vBNS-NSF Relationship

"The relationship between the RA, vBNS and the Government is essentially unworkable given the present limitations the Foundation operates with. By splitting the NSFnet Backbone award into two separate pieces, the Foundation puts itself in the role of a systems integrator, which it cannot effectively perform without a dedicated Government technical staff to perform supervision. The Foundation cannot tell the RA and the vBNS to go off and work together. It cannot tell the vBNS supplier to do whatever the RA directs, and neither can it tell the RA to act as the vBNS provider directs. When problems develop, there will also be a temptation for the RA and vBNS providers to blame each other, leaving the Government in the middle to try and arbitrate."

"The relationship between the RA, vBNS and the Government is essentially unworkable given the present limitations the Foundation operates with. By splitting the NSFnet Backbone award into two separate pieces, the Foundation puts itself in the role of a systems integrator, which it cannot effectively perform without a dedicated Government technical staff to perform supervision. The Foundation cannot tell the RA and the vBNS to go off and work together. It cannot tell the vBNS supplier to do whatever the RA directs, and neither can it tell the RA to act as the vBNS provider directs. When problems develop, there will also be a temptation for the RA and vBNS providers to blame each other, leaving the Government in the middle to try and arbitrate."

"The Foundation's track record in terms of technical management of the existing awardee is less than stellar. This is not because the Foundation's personnel are incompetent or unwilling to perform supervision; it is because the existing staff is hopelessly overburdened with other work, and simply does not have the time or resources to perform adequate supervision. If this arrangement were effected by NASA or DOE, it might be doable, because they have large in house staffs that can support the type of management needed, but this structure adds even more work for the NSF staff to perform. Thus it is imperative that NSF structure the solicitation in a fashion that it is capable of managing effectively."

"The reason that is often given for splitting the RA and the vBNS provider roles is that the vBNS provider has a bias towards its own network, and cannot be counted upon to act impartially with respect to other providers. There is an assumption here that the vBNS has other, non-NSF supported traffic that it carries over the same facilities and routers that are used to provision NSFnet backbone services. If it had no such traffic, and were exclusively dedicated to servicing the NSFnet backbone community, then it would not have an opportunity to show favoritism. This would obviate the need for splitting up the backbone services award into two segments, and allow a single organization to provide operational management and support of both routing control and network transport, preserving the end to end management role that the existing awardee has, and which has served the Internet community very well over the years."

"No organization which has experience operating and managing a national backbone would separate routing management from network transport if operational stability were the pre-eminent goal. The advantages of having a single focus for both management and responsibility of the overall system are very significant.. . . . From the Government's side, this approach results in one organization that is responsible for meeting milestones and performance goals, and provides a single point of contact for new service introduction and interface with other agencies and networks."

"No organization which has experience operating and managing a national backbone would separate routing management from network transport if operational stability were the pre-eminent goal. The advantages of having a single focus for both management and responsibility of the overall system are very significant. This is even more important as new technologies and routing mechanisms are evolved into the network. From the Government's side, this approach results in one organization that is responsible for meeting milestones and performance goals, and provides a single point of contact for new service introduction and interface with other agencies and networks."

"NASA strongly recommends that NSF seriously re-examine the rationale of splitting these services amongst an RA and a vBNS provider in the light of operational stability as the pre-eminent objective. NASA also recommends that NSF perform an internal management study that outlines the resources required to manage a single awardee with those required to manage both the RA and the vBNS provider in the roles outlined in the solicitation, and should definitely include staff from grants and contracts in this study."

Cost Issues

". . . . Splitting the award into two pieces will be more expensive than having a single integrated award. As outlined before, both the Routing Authority and the VBNS provider must have 24x7 staffed operations and management organizations. This will cost more than a single integrated operations and management staff. The travel and coordination requirements between the NSF, the RA, and the vBNS provider will also cost more than if only the NSF and an integrated provider were involved"

". . . . Splitting the award into two pieces will be more expensive than having a single integrated award. As outlined before, both the Routing Authority and the VBNS provider must have 24x7 staffed operations and management organizations. This will cost more than a single integrated operations and management staff.

"There is also an argument that the

NSF cannot afford a dedicated high speed backbone and thus that these backbone facilities must support shared use, and therefore there will be a built-in bias by the network provider against other network organizations and so the routing and transport organizations must be split apart to preserve fairness. This is probably true if one looks at a network architecture built of IP routers with point to point DS-3 trunks in the manner the existing awardee has constructed them."

"However, the NSF is clearly trying to push in the direction of Broadband Integrated Services Digital Network by specifying 155 Mbps transmission rates for the vBNS provider. Given a B-ISDN substrate as a base, the core network switches would not be IP routers at all but something more like an SMDS or ATM switch based infrastructure, which could also be used even if the requirements were only at the 45 Mbps level. In this architecture, the IP routers only act as interface points from the site networks to the BISDN cloud. Thus the relatively inexpensive routers could be dedicated while the switching fabric itself could be shared. Different levels of service could be provided to each access point according to need. in this architecture the cheap parts (the routers) are dedicated, but the expensive parts (the high speed cell switches and network transmission facilities) are shared."

[Editor - this ability to buy into a shared substrate of the public switched telephone network would be lost should the NSF accept IBMs or ANS's likely bid that will focus on proprietary Planet, Orbit and Comet switch, LAN and router technology. It seems that such a move would clearly run counter to Congressional intent that technologies in NREN conform to telephony standards.]

Given a B-ISDN substrate as a base, the core network switches would not be IP routers at all but something more like an SMDS or ATM switch based infrastructure, which could also be used even if the requirements were only at the 45 Mbps level. In this architecture, the IP routers only act as interface points from the site networks to the BISDN cloud. Thus the relatively inexpensive routers could be dedicated while the switching fabric itself could be shared.

"If proper performance criteria were established, a 'dedicated' (at the IP and CLNP level) network could be constructed, and proper levels of service assured, without having to deal with the issues involved in a shared IP or CLNP infrastructure. This is the very approach that DOE and NASA are taking in constructing their HPCC networks. These networks certainly would not tolerate sharing of their router facilities and policy controls [with commercial carriers] but they are more than willing to share transmission infrastructure given proper performance assurances."

"This approach yields a relatively low cost for high performance network facilities, but avoids the issues of actual and perceived bias of the control of network layer entities, such as routers. Applying this approach to the NSFnet backbone recompetition would allow an integrated network operations and management structure, and still achieve high performance at acceptable cost. It all depends on what is shared and what isn't. Fortunately, in NSF's case, the part that needs to be dedicated, i.e.; the routers, is relatively inexpensive and have low recurring costs to the program."

"This approach yields a relatively low cost for high performance network facilities, but avoids the issues of actual and perceived bias of the control of network layer entities, such as routers. Applying this approach to the NSFnet backbone re-competition would allow an integrated network operations and management structure, and still achieve high performance at acceptable cost. It all depends on what is shared and what isn't. Fortunately, in NSF's case, the part that needs to be dedicated, i.e.; the routers, is relatively inexpensive and have low recurring costs to the program."

"Also, there is no statement in the solicitation that the vBNS must be committed to offering the proposed network service as a commercial service, available to other agencies and private industry, which is a goal of the NREN initiative. By trying to push the vBNS service in a direction which is consistent with BISDN offerings, NSF would accelerate the deployment of these services in a commercially viable manner, resulting in lower overall cost to the Government."

"By not stressing the desire to use Commercial Off The Shelf (COTS) products, NSF will only encourage the development of special purpose, one of a kind hardware and software, and prevent the amortization of development cost over a much broader community of users, and also incur higher system maintenance costs. Including a statement that strongly discourages the use of one of a kind non-COTS oriented technology in the solicitation would lower costs to the Government,

"By not stressing the desire to use Commercial Off The Shelf (COTS) products, NSF will only encourage the development of special purpose, one of a kind hardware and software, and prevent the amortization of development cost over a much broader community of users, and also incur higher system maintenance costs. Including a statement that strongly discourages the use of one of a kind non-COTS oriented technology in the solicitation would lower costs to the Government, and increase the technology transfer aspects of this NREN project to the network community as a whole." It will be VERY interesting to see whether the NSF takes NASA up on this suggestion. Doing so would close the door to the very eager attempts of IBM to insert its proprietary Planet gigabit switch, Orbit gigabit LAN, and Comet router.

"NASA recommends that NSF re-examine the solicitation's structure from a cost perspective, given the aforementioned B-ISDN technology options that NASA and DOE are both pursuing at this time, and should be relatively mature at the time of the transition to the new award. It may be possible to achieve lower costs through an integrated organization, and yet still achieve high network reliability and performance at the same time. Even more saving and higher productivity

could be achieved by stressing the importance of the vendor using COTS technology wherever and whenever possible, and thereby also advancing the NREN objectives of commercialization of technology and the development of high performance network infrastructure."

Program Priorities and Objectives

"The solicitation outlines a number of objectives and requirements for the NSFnet backbone. However, it does not specifically outline the relative priorities and objectives of the NSFnet NREN program. Such a statement of priorities would help the vendors overall understanding of the solicitation, and what areas are more important than others. This would also help in clarifying actions if budget cuts were to be absorbed by DNCRI's networking program. This will be a likely occurrence in the current budget climate in Washington. A credible internal plan here, would help both the awardee and the Government in it's ability to flexibly restructure the program in a changing budget climate."

"NASA recommends that NSF add a section to the solicitation that outlines NSF's overall NREN vision and relative priorities of the various objective and requirements present in the solicitation."

Technical Review Panel

"In light of technical complexity of the network design, and new service evolution, substantial technical review of the vendors will be essential for successful transition from the existing network system to the new one. However, as was mentioned before, there is no staff support inside NSF that can adequately provide this sort of technical review. Therefore, we suggest that an independent NSFnet backbone technical review panel be formed, and meet on a quarterly basis, or when significant technical decisions are about to be made. This panel should be comprised of experts in the fields of routing, network management, operational support, network transmission and carrier systems, and general router design, which would have access to proprietary vendor information on a non-disclosure basis."

"In light of technical complexity of the network design, and new service evolution, substantial technical review of the vendors will be essential for successful transition from the existing network system to the new one. However, as was mentioned before, there is no staff support inside NSF that can adequately provide this sort of technical review. Therefore, we suggest that an independent NSFnet backbone technical review panel be formed, and meet on a quarterly basis, or when significant technical decisions are about to be made.

"The panel members would come from other agencies, academe, and the commercial sector. They would not be allowed to make policy decisions of course, but would make recommendations to the Government about how the vendor or vendors is proceeding in areas of engineering, transition planning, deployment and operations. They would also provide feedback on new directions and overall network evolution. They would likely have to come from organizations which had no conflict of interest with the awardee however. These would not be network users or administrative types, but people affectionately known as network 'weenies' or 'wizards.'"

"Had such a panel existed during the DS-3 upgrade design and deployment, significant amounts of problems may have been avoided, and certainly the NSF would have been perceived as being more in control. *The NSFnet backbone should really be viewed as a national asset; the centerpiece of the national network infrastructure.* [Italics editor's] Such a panel would add credibility to that viewpoint, and provide the Foundation with much needed technical feedback, which is critical for good program management."

"NASA recommends that the Foundation form an NSFnet technical review panel, composed of experts in their fields, whose purpose would be to review the vendor's technical approaches and designs and comment on transition and evolution strategies. . . ." As we say in our editorial, the NSF stands before a clear fork in the road. We hope it makes wise choices.

Subscriber Survey

How does the COOK Report meet you needs? Please respond to these questions and add any comments of you own. Mail or fax (609 882-2572) this page to me. If faxing please call first so I can turn on my machine.

What I like best about the COOK Report is: _____

What I like least is: _____

Too much coverage of: _____

Too little coverage of: _____

Layout -- would you prefer larger type and more whitespace? _____ Would you accept an increase in your subscription cost of between 15 and 25% to pay for increased printing and postal costs?_____

Other comments _____

_____

_____

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