The COOK Report on Internet

SURFnet Users Pay Cost of Connection While Gov't Funding Buys Advanced Infrastructure Capabilities That Serve as Enablers of New Applications

Like Canada an Emphasis on Dark Fiber Pushes Network Forward While Other European Nations Focus Much More on Applications

Volume X, No. 1, April 2001
ISSN 1071 - 6327

Editor's Note: Kees **Neggers** is Managing Director of SURFnet, the national research and education network of The Netherlands. He has been involved in research networking since the mid-80s when The Netherlands started to build its national research network. He is a trustee of the Internet Society. He has a 1972 degree in electrical engineering from the University of Eindhoven. As he explained it: Informatics and computer science were not really available at that time. You either had to do mathematics or electrical engineering. We interviewed Kees in Montreal on November 29, 2000.

COOK Report: I gather that you are building the next generation of SURFnet, a 10 gigabit DWDM based network. Where does what you are doing now fit into the bigger future?

Neggers: SURFnet's mission is to provide advanced network services to the research and higher education community in The Netherlands. As such we are always pushing the edge. In fact, as soon as we would not do this anymore we would become redundant and our users could simply buy services from the open market.

COOK Report: Now, who actually owns and possesses the fiber that you will use?

Neggers: At the moment KPN Telecom does. KPN Telecom is the provider of SURFnet4, our current operational network. KPN Telecom is the private successor of the national PTT and has been an excellent partner in setting up the SURFnet organization since the day we started operation in the second half of the 80's and with the creation of SURFnet4 in particular. The tender for SURFnet5, our next generation network was won by Telfort and Cisco Systems. Telfort also will own the fiber. They will give us lambdas on that fiber.

COOK Report: And when you say 'tender', you mean contract?

Neggers: A tender means an open procurement according to the rules of the European Commission.. That's the only legal way government money can be spent in Europe. We received a substantial subsidy from the government for our networking activities. This is the only way to stay ahead of the commercial networks. In the past, we didn't need to tender, because there were no alternative providers, but this is now an open market and for everything you acquire in the open market, you must make a tender. The nature of the resulting contract by the way is more partnership for a common development than just a normal supply contract. Research people from both Telfort(BT) and Cisco Systems are actively involved in our activities.

COOK Report: So you can really just bid, taking the funding that you have, for lambdas or a whole strand of fiber?

User Income Builds Each Successor Network

Neggers: In establishing SURFnet, in 1986-87, the government in the Netherlands took a deliberate decision that what it created should both be a company in order that it could make decisions quickly and be managed by the user community with the assistance of the PTT. While the government promised to pay for the new capabilities of the network, the users right from the beginning had to pay for running the network. This model has been useful ever since. We have always had the income from the users to maintain the network and keep it operational. Moreover we always got money from the government, which we only could use to push the envelop of what the network could do. The users are tempted to spend their money on things — like traveling and ordinary equipment — things other than innovation. So in the end, this money saved us.

COOK Report: Was this distinction different from many of the other European models?

Neggers: Yes, different from all of them, al-

On the Inside:

though several have introduced this model in some form as well now.

COOK Report: A mutual friend of ours dislikes Internet 2 or Abilene because of what he calls the Chief Information Officer's syndrome. He finds that the university CIOs want the government money and subsidies just to pay for the cost of the networking connection. You're contrasting your situation to this type of mentality?

Neggers: Yes. The other fairly important difference between Abilene and Canarie and SURFnet is that we provide full global Internet connectivity to our users. That means we are their ISP. And that means that our customers, the universities, rely on our service for everything they do.

COOK Report: It gives them a very strong interest in making sure everything is working well?

Neggers: Yes, because it's their lifeline to the Internet. But it's also interesting for the carriers and equipment providers to work with us, because we have high standards to maintain. We are not just experimenting.

COOK Report: I like it. Your situation allows you to cross the boundary between a test bed and a production network.

Neggers: Yes. Here is how we did it. From day one, with the "innovation" money we built a network — SURFnet1. With the user money, we made it operational and kept it operational. In parallel, we used the government money to build the next generation network. Now we are building SURFnet5 with government money and the users are paying for SURFnet4 which is essential, because it is the lifeline for our customers. Because we have to be as advanced as possible, we have to take risks in our procurement. We set new standards. But we cannot introduce the universities to them before they are proven technology. So, while universities at the moment are connected to SURFnet4, they will have a secondary connection to SURFnet5. We will use their production traffic on SURFnet 5. If SURFnet5 fails, they can fall back on SURFnet4. And as soon as it's stable enough, and the Quality of Service is in line with our expectations, we will turn off the old network.

COOK Report: So you're putting SURFnet5 together how? Out of what kind of components?

Catalyzing Change

Neggers: Two components. Ten gigabit lambdas from Telfort. And GSR+ routers from Cisco Systems. That's all.

COOK Report: What is Telfort?

Neggers: Telfort is the 100% subsidiary of British Telecom in The Netherlands.

COOK Report: So if you have ten gigabit lambdas, is it almost the same thing as having your own fiber? You have the pipe that brings you the bits and you can decide what to put on the end of it, how to structure it, and how you connect it?

Neggers: Yes. We don't own the fibers, but we do fully control what we do with the lambda's. We operate the network. We build the network out of these two components. We have our own subcontractor, which is the SARA computing center in Amsterdam, to operate the network for us. So in effect, the advantage of this solution is that you can share the cost of the fiber. The result is now SURFnet5. When we made our tender for this, we really pushed the edge of the infrastructure and networking, and not just in The Netherlands, but also in Europe and even beyond that as well. Because what we did is to ask for things that we think are technically possible, but for whatever reason the commercial market is not yet willing to deliver.

By having this government money, real money, and by saying this should be possible, we just put out this procurement and invited the market to react to it. And we got five offers for our desired network. And if we had not tendered for SURFnet5, all these people would still be trying to sell us ATM and SDH. By leading them with our tender we educated them.

COOK Report: Pulling them up by their own bootstraps.

Neggers: Absolutely. And not only this, but also by having these fifteen points of presence being supplied by Telfort we caused them to have to bring fiber in which led to the realization that others could bring fiber in along with them. If someone asks to lay fibers, the city has to coordinate it and everybody has the opportunity to use the same trenching of the street to add their fibers. That's why you see all these different colors at the same time in the photographs in my presentation.

COOK Report: Are you looking at what Stokab does in Stockholm?

Neggers: No, we don't have that model in The Netherlands yet and, obviously, in Amsterdam you don't need it, because the supply market is already doing it. But for these fifteen concentrator locations it is a different matter entirely. They are spread all over the country. Because Telfort had to bring in dark fiber, all the others added their own fiber and there was an investment of infrastructure where none existed before. The city governments have the obligation to coordinate fiber builds. Because of this every other business in The Netherlands has the opportunity to join anyone who is laying fiber, since local governments don't want you to tear up the streets every month.

COOK Report: Let me ask you what may be a political question. At the presentation this morning, the Dante Geant project map of Europe was conspicuous with The Netherlands' absence. Putting two and two together, I would infer that there's probably a different philosophy going on there. Is that a philosophy you feel comfortable saying anything about?

Innovation by Committee Doesn't Work

Neggers: I have basically two comments. First SURFnet has found that European networking so far tends to be driven by the speed of the slowest. I don't like that model. It is a recipe for not being able to be state of the art. We wanted to use the opportunity of the setting up of Geant to learn from the past and improve the situation. However the way the Geant network is organized is a continuation of the structure from the past. Dante has no central management. It's a consortium of some twenty-six national research networking entities. All 26 have to agree on everything. I didn't want to be the 27th of that group. So my proposal was, Dante should do the procurement, should operate the network and I want to be a customer of DANTE. And the consortium should only be a consortium in its relationship to the European Commission to define the set-up. But none of the participants in Europe were willing to go that route.

The second objection I had is that we have in Europe an association of national research and education networks (NRENs), called TERENA., DANTE was created, with the help of TERENA to provide operational network services to allow TERENA to concentrate on the development activities. The Geant network should not be a development project, it should be delivering services. In the past, these Commission contracts were always affected by research demands and all the operational networks were in effect research projects. This time, for the first

time, the commission said explicitly that it no longer wanted it that way. They learned from that. But DANTE and the other NRENs still seem to like this model and want to continue playing around. These two points of view were the main reasons that we didn't want to participate in the Geant Consortium. In this, SURFnet is fully supported by the government, because this has already been escalated through the proper political layers, etc. Connectivity is of course a different story. We are well connected to the existing TEN-155 network and we plan to connect to the Geant network too.

COOK Report: I like independent thinkers.

Neggers: Remember that SURFnet was also the organization who started the EBONE initiative. This really changed networking in Europe. This was 1991 and at that moment in time, all politicians in Europe wanted OSI to be a success. They never wondered how these protocols could ever get to work, and who should be expected to deliver the equipment to do it. We still had to sit and wait for the industry and the telecom players to solve this.

COOK Report: Well, OSI was done by a committee and so was DANTE. It seems that the TCP/IP philosophy of does it work in the real world or not is more congenial to your point of view?

Neggers: Yes. And, as a result, SURFnet is now connected to the current TEN-155 network, with 622 megabits per second via gigabit Ethernet interface. We are the only NREN in Europe with a 622 megabit connection to TEN-155. The next best are the 155 megabit connections. From the very beginning, no one has been able to upgrade until now.

COOK Report: I'll bet it hasn't happened because of the committee process environment that runs it? You have to have so many different agreements from different groups and users and providers?

Neggers: I'll give you a typical example. The TEN-155 core ring is now a 622 megabits ring. And SURFnet is currently the only NREN connecting at this full bandwidth. And this ring goes Amsterdam, Frankfurt, Paris, London, Brussels, Amsterdam. So five countries. Only Germany and the Netherlands have committed to have 622 access. This ring goes through London, Paris and Brussels and they are not using the opportunity to upgrade to 622. I don't know why. This is a waste of money and resources, isn't it?

COOK Report: It makes sense. The end user bandwidth you get from the Telfort fiber is gigabit Ethernet. Did I hear that right?

Neggers: No, the architecture of the network is as follows. Telfort has their national backbone network that includes their fibers, with dense wave division multiplexing equipment. And this touches the Telfort points of presence. From Telfort points of presence, there is, by two diverse routes, dark fiber to our points of presence. And Telfort delivers the lambda's over this dark fiber to our Cisco GSR+ routers on all our points of presence. So there are no switches in-between.

Gigabit Ethernet

COOK Report: But then you have your own network, right? And as you distribute from your own headquarters.

Neggers: To our customers. This is what I call the last kilometer problem. There I also need dark fiber, or gigabit Ethernet at the moment. This is going well in several locations now. And, of course, we want it everywhere.

COOK Report: Then you have the ten gigabit lambda coming in and you either have put or will soon put gigabit Ethernet on it by means of switches?

Neggers: We will put gigabit Ethernet interfaces in our GSR+ routers to connect to our customers.

COOK Report: Like an interface card for gigabit Ethernet port?

Neggers: Yes. You have now cards with three gigabit Ethernet ports on it.

COOK Report: And what are the costs of those cards? Do you know?

Neggers: I don't know these prices by heart, but the gigabit Ethernet is substantially cheaper than the SDH cards. We are aggressively changing away from SONET and this is going on everywhere. So dark fiber with gigabit Ethernet on it is much cheaper than SDH. No matter whether you do it yourself or buy it from a provider, this is the main reason that we are switching so fast. It's also easy to manage.

COOK Report: Therefore you use gigabit Ethernet streams or ports for the customers whom you can reach with current fiber attachments, and the ones you can't, you're going to be laying dark fiber to.

Neggers: Our advice to everybody is to own your own dark fiber to a neutral exchange point where you can meet all kinds of service providers. That's our model of networking. And that's what we stimulate in the region where we have our customers. In our next generation networking, we are no longer just talking to carriers to make plans, we are talking to customers to make plans. And the carriers are welcome to deliver what we ask. But you see, we have the strengths of the end user community to tell them what we want. As with SURFnet5, we told them what we wanted : 10 gig lambdas, rather than ATM and SDH.

COOK Report: That's a major seismic shift.

Neggers: We get government money to play this role, because the government in The Netherlands is aware that information has to flow over fibers. It knows that the sooner we can develop widespread fiber-based infrastructure in the Netherlands, the sooner the whole population, not only in Amsterdam but in every remote corner, can participate in the next economy.

The current set-up is two separate ten gig lambdas that we drive with packet over SONET framing. The only thing I can offer is an Internet service. For the next generation, optical switches and lambda switching should be an option as well and then you can build powerful grids. The only way to build grids at the moment is via virtual private networks, which we will do as well, of course. Our backbone is still powerful enough. So, all the grid people in the Netherlands I talked to said: "Don't worry. For the foreseeable future, your capacity is good enough to use grid networking as virtual networks."

Grids: Prospects and Dangers

COOK Report: How would you define what grid networking is all about and the concepts?

Neggers: Grid networking is nothing else than a closed user group building a private Internet among themselves, which still is then connected to the wider Internet, either via the institutions or via this common network. But in the old days, you had a physics network and a space agency network. Each discipline had their own network. In the future, this will become possible again. There are some advantages if you are in a closed network. Until now, the disadvantages of having to buy your own wide area lines in order to do your own complex management has discouraged all these user communities

from doing it. But now you see it coming back, because wide area networking is becoming cheaper. Furthermore by using gigabit Ethernet, at distances up to 100 kilometers, you can have quite powerful networks that you manage like your own personal Local Area Network (LAN).

COOK Report: And even further than 100 kilometers now.

Neggers: Yes, so it's changing. So the next generation networking, SURFnet6 I think will need dark fibers to provide multiple lambdas for our customers. We need lambda switching to do that correctly. We also need to integrate that with the Internet and this is exactly what we want to start with the lambdas that we try to get across the ocean, with the Chicago and the CANARIE people.

COOK Report: How is this related to OBGP?

Neggers: Well, OBGP could be the protocol to merge lambda switching and Internet, because what we don't want is lambda switching to become a circuit-switched network. Because circuit-switched networks is the wrong concept in networking. We will use MPLS in SURFnet5 for two reasons. First, to do the fast reroutes. If one lambda fails, the router has to detect it in milliseconds and only use the other one, because we don't have SDH to give this protection.

COOK Report: But Van Jacobson and two other people at NANOG 20 gave a paper on what has to be done to get to millisecond time in recalculating routes.

Neggers: SDH gives you 50 milliseconds or less and we are very confident that Cisco will be adequate. When we started EBONE, we were the first wide area network to deploy BGP3 and allowing us to work with Autonomous Systems for the first time. So the next generation BGP was made operational by EBONE engineers like Peter Lothberg, Bernard Stockman and Erik-Jan Bos. Sometimes you need to change things to be able to take the next step.

COOK Report: I think it becomes useful then to see the type of thinking that you're building your network on seeded into other countries, so you can talk to each other with your same resources at the same level.

Neggers: The last thing SURFnet wants is to stay ahead too much of the rest. We want the rest to go as fast as technology allows too. And only if everybody will do that we will be able to talk as fast as possible among each other.

COOK Report: And do so into lots of different places.

Neggers: As I said earlier, the grid thing is just a new opportunity for communities, like high energy physicists, astronomers, whatever, to build their private network. And have it dedicated for them and use it tuned to their needs. No general purpose network will suit all specialist needs well. Although a good Internet is, of course, as good as a private network. So the better the network there is, the less need there is for a dedicated network.

But with the availability of lambdas all over comes a challenge to build your own networks. So as I said earlier, in the Netherlands, currently, all the grid people just want the connection to SURFnet5 and that will satisfy their grid needs for a while. The next SURFnet might well be optical switches with lambdas, where in fact institutions sit on a number of lambdas themselves and tune the bandwidth to the rest of the world based on their needs. And there they need OBGP to do that in a scalable way. Because we have no interest in a fragmented Internet, that would be made up of numerous amounts of small grids, with nothing to connect to. In particular, if the grids would go in diverse directions with different protocols, we would in fact destroy the Internet, rather than improve it. That's what we need to be careful about.

COOK Report: What would be the danger of going in diverse directions? People might make protocols for physics and some other people might make protocols for medical specialties or whatever?

Neggers: Yes. And it simply doesn't scale. For instance, in The Netherlands, the high energy physics community had their own network before SURFnet started. However they were more than happy to start using SURFnet once we'd proven to be a reliable service provider and they have benefited from using us ever since then. They are not driving towards their own solutions, so we still work very closely together with them and the first extension of an Amsterdam optical exchange connected to Chicago might well be to CERN as a corporation of The Netherlands physics community and the CERN physics community.

COOK Report: The CERN in Geneva?

Neggers: Yes. In fact, Amsterdam and CERN have very often been the two hot spots in Europe for the next generation networking.

COOK Report: Are you referring to the Web, with CERN?

Neggers: No, the Web is a completely independent development that took place at CERN. But, for instance, the first private X.25 network in Europe was a link between CERN and Amsterdam. And we made that into a 256K link by joining forces of a EUnet link, one of our links, an IBM link for EARN etc. And also, the two main points to start EBONE were CERN and Amsterdam and then we added two more to make it into a carre.

OBGP and Issues of Network Architecture

COOK Report: You seem to be well aware now of OBGP and how it can be useful. I gather Tom Defanti has also your level of awareness. How have you all become more aware of it? Has Bill been talking about it in meetings that you've been at over the last six months?

Neggers: My source for OBGP is Bill St Arnaud.

COOK Report: He has that pretty long paper that he shared with me at the beginning of the summer, when I asked about it.

Neggers: He spoke at length about OBGP at the NORDUnet conference in Helsinki at the end of September.

COOK Report: Was this your first introduction to it or had you run into it earlier?

Neggers: OBGP? It was the first time that I was exposed to more details of it. I'm not doing the technical work myself. But my question directed to lambda switching was always do we want to move to a circuit-switched network again? And the OBGP was the first time people told me, no, there's no need to do that, it can be done differently. I personally don't see how the circuit-switched networks can fit into the Internet. What can you do with it, other than make point-to-point connections?

COOK Report: And, eventually, when you've got point-to-point meshes numbering in the hundreds of thousands and beyond, it gets a little complex to manage.

Neggers: Yes. Also I'm a strong believer, like Bill St Arnaud that the network should be stupid and reliable and the edges should be smart. That scales. Architecturally the circuit-switched networks are the other way around.

COOK Report: What are the implications

for the backbones? I've heard people like Sean Doran, for example, express concern about UUNET's backbones that tends to be made up of, what? Thousands of flows, circuit-switched flows? Are you thinking that we're pushing the limits of scalability of that type of network architecture right now?

Neggers: I have no particular knowledge of their limitations.

COOK Report: But if you're doing predominantly circuit-switching with ATM or whatever as opposed to the more traditional, connectionless networks, my question was does this have implications for the scalability of backbones at that level? Is that part of the problem or not?

Neggers: I think a backbone provider could manually, by an operations center, change the topology of its network. That could be by a circuit-switched solution, to add lines and change lines. But such an action is only an effort to optimize the backbone. It is nothing more than traffic optimization, or traffic engineering. It won't scale in the end. However, with OBGP, each network can decide to take another route to another network, if there is enough demand.

COOK Report: If you look at your backbones now that you're running to connect the international traffic to DANTE to Geant or wherever, what is the situation now that you are faced with in scaling these backbones upward as you get more and more traffic? If you get traffic from Netherlands's universities and users in general that asks to go other places in the world, either the traffic has to go there via a relatively small number of large backbones or by the more diverse OBGP kinds of routes, right? And is part of the motivation for going in the OBGP direction the concern about the continuing scalability of your backbone connections that you need to interconnect with the rest of the world if you don't have these things?

Neggers: In fact, BGP is doing this right now. We have to manage our peerings. And at the moment, we do that at the level of choosing this upstream provider or that one. So we have to decide, in fact, the traffic has to go to Teleglobe, Level 3, DANTE or the Amsterdam Internet Exchange.

COOK Report: That is for your upstreams?

Neggers: Yes. And the Amsterdam Internet Exchange is the most important because it has many peerings. So in fact you decide which network you peer with and you direct the traffic towards these peers and then you need a last resource upstream provider and you dump the rest to that provider as a transit network.

Here we have two choices, Level 3 and Teleglobe. We split the loads that we send each of them based on what we think is the best network to carry to certain destinations. So we have to do that management work as a provider for our customers. With OBGP we are not only able to do that on a, let's say, OC-3 or OC-12 or OC-48 connection, but we will be able to do that choosing dedicated lambdas to certain destinations. So if we see that there is a large demand to certain destinations, you can then add shortcuts via lambda switching to them.

COOK Report: But what you say implies that there are some upstream providers with an entire lambda of ten gigabits to spare?

Neggers: Not yet, but this is a likely development.

COOK Report: Then it doesn't make sense to talk about less than 10 gigabits for a lambda? Are you aware of anyone talking to upstream sources of fiber for entire lambdas at this point that might be usable in this way?

Neggers: This is pure research, Gordon, because we don't have the switches and protocols to do that at the moment. The dark fiber for it is becoming available, many places. The lambdas, the dense wave division multiplexing equipment is becoming cheaper and stronger, but we don't have the switching devices.

COOK Report: I have heard that you need very, very simple switches, much more simple than all this stuff that's on the market right now.

Neggers: Yes. But I also want to be able to manage the switches via the same management tools that I have to manage my BGP peerings.

COOK Report: That's why it's desirable to use BGP in the switching process?

Neggers: Yes.

COOK Report: Marc Blanchet who is really the co-developer of OBGP with Bill St. Arnaud was telling me he doesn't see any roadblocks that'll make it impossible now. However he adds that it hasn't gone far enough to be absolutely sure that it's going to be feasible with BGP. But he also says that the need is presumably going to be enough that, if for some reason you can't use BGP, somebody's going to have to probably write something that will do what is wanted. [**Editor's Note**: as we have since learned, as a fail safe device, St Arnaud and Blanchet have a team at Carleton University in Ottawa that is doing this right now.]

Neggers: That will be a smart idea on how to overcome that problem. Because the challenge to find the solution is definitely there.

COOK Report: With all the conduits containing more than 400 strands of fiber that are now going in the ground, you have to think that not only are you doing yourselves a favor with OBGP, but hopefully, you're doing some of the fiber owners a favor.

Empowering your Customers

Neggers: Yes. The reason that Bill St. Arnaud and I are interested in this is that we are not providers. We work for a user community. If it is better for the user community to get many lambdas on their premises, we will deliver them. However, a provider might well want to keep a provider relationship where the provider is the smart network provider and it can force you to work through it.

COOK Report: There's a number of advantages for your customers in being able to put their switching hardware on, attach it to the fiber in such a way that they can determine the use of their lambdas.

Neggers: Well, the capacity of these lambdas is enormous, of course, and at the moment, we don't have any devices by which to fill them, other than by very expensive routers . No single end station can use this much capacity. But in the end, of course, local area networks will have special facilities that demand special needs, not all of the time, but some of the time, to certain destinations.

COOK Report: So there should be some, perhaps, switching center in Amsterdam that can talk to the OBGP confederations of the world and show your customers, almost like an airline scheduling board, that lambda availability for the next 24 hours, 7 days, 30 days, from these various points to those points is thus and such. And if you want to order a lambda, tell us what you want?

Neggers: The interesting thing about OBGP is that you don't ask a lambda for an end station to an end station, but that you must deal with this at the institution level. So the institution will then have a peer-to-peer connection with another institution or another provider and then everybody can use that connection, not just the single end station that wants it; it is part of the infrastructure

after that. That is the whole trick, in fact.

COOK Report: Yes, and the University of Amsterdam or Princeton University or Stanford University, the entire university could fill a lambda much more easily than any one laboratory or one application

Neggers: Yes.

COOK Report: And I suppose it's pretty easy to imagine a big, major university being able to use a whole lambda in a year or so.

Neggers: In particular because it's obvious that the Internet is going to be used for everything. It's running on top of everything and it will be used for everything now. You still have this Quality of Service problem. By having these lambdas to come in when you need more bandwidth, you overcome that completely. For most of the applications, if you have the ability to sufficiently overprovision your network, your Quality of Service issues just disappear.

COOK Report: As Brian Reid said in talking about Quality of Service with regard to the community network being built in Palo Alto with a gigabit to every house then who needs it? Only telephone companies think you do.

Neggers: Well, currently, we run high quality videoconferencing over our trans Atlantic links. We do this without Quality of Service features and it works fine. You can protect some of these streams, but it is also a difficult process. You can't direct from a management center the quality of all the streams by giving them special treatment, by excess bit rates or whatever you want to attempt. It's just not scalable.

And at the moment, we are in the position that the wide area network is not the bottleneck. In the initial days of networking, wide area network was always the bottleneck. And from SURFnet4 onwards, this disappeared. Now, the local area networks are the bottleneck, or even the end systems are the bottleneck for the communication. Gigabit Ethernet is becoming a commodity. If you buy an Apple computer, you now will get the gigabit Ethernet. This will change the world. People will find solutions to use that.

COOK Report: If I'm Stanford University or if I'm one of your universities and it's possible and affordable for me to buy a lambda, are you really restricted to buying it between universities or between network operations centers like your network? In other words if Stanford buys a lambda, then maybe they can let everybody, all their labs and so on use it, but can they really? Because doesn't it imply that the users of it all want to go to the same place. Or if they don't want to go to the same place, then how do you transit it to the rest of the Internet? And, again, what I was thinking is if the universities can have these huge pipes, not only on their own campuses, but their own departments, but they can to some extent become bit providers to business around them, and to laboratories around them. Do you see what I'm thinking? Is that feasible or not?

Neggers: Well, you're entering now into a completely different area of problems, that is the anti-competition stuff. In Europe, at least, anti-competition regulation is very stringent and a subsidized university should not compete with commercial operators. The same holds for SURFnet, so we provide service to a target user group and there is no reselling of our services in competition with commercial operators. That would be completely illegal, it would kill our partnership with the government.

COOK Report: I certainly understand. But supposing you had a community fiber network like Brian Reid was talking about in Palo Alto. Brian made it very clear that it's a Level 2 network only and it's going to become open to anyone who wants to use its infrastructure to provide Level 3 services and beyond. What I'm wondering is a community network like that with Level 2 infrastructure, possibly transferable? You know the situation in Europe, if you imagine this existing in Utrecht or Rotterdam, or Amsterdam, wherever, or even some smaller town, what would be the situation there? Would it be different, as far as the competition goes? If the university wanted to take a surplus part of its lambda and make the bandwidth available to a community network that was only layer 2 and up. I suppose it wouldn't be any different, because then the commercial carriers would say you're bringing them higher level transport at a cheaper price than we can do it.

Neggers: It would kill the whole thing. So if you are funded by government funds, you always must stay out of competition with commercials. And, in fact, if you spend the government funds, you have to tender for that. You say this is what I want and then the market can provide it. If the market is not responsive to tenders, you are free to do what you want.

For the projects that we do, one of the requirements that we put in is that a dormitory that is housing for students and non-students, only the students get access to SURFnet. And they use smart cards to do the authentication and the non-students belong to a commercial Internet provider and the others go to us. And this is part of our tender and this solution will become available for everyone who wants to deploy it elsewhere.

COOK Report: Isn't it likely that, once OBGP is developed in these Canarie and SURFnet communities, companies which are primarily owners of raw commercial fiber rather than carriers — and I'm thinking of Level 3 and Williams and PFNet and Aerie and some of the new ones, as opposed to Sprint and AT&T and MCI — would find setting up their own infrastructure using these techniques desirable?

Neggers: Absolutely. That's what we're doing with nearly everything we do. If we see things that are technically possible and useful, we force them into existence.

COOK Report: And when you do that, somebody like Level 3 and Williams, who are wondering what to do with all their strands of fiber have to watch it and say, Oh, maybe we should get together and set up our own networks on this basis and compete with the older PSTN carriers on a commercial basis. Is this the best way to make cheap bandwidth available to business and industry?

Stimulating Startups

Neggers: In fact, you have to stimulate the start-ups. Incumbents have vested interests and sometimes it's not their primary interest to move fast and use new solutions. So you have to circumvent them via smaller, new companies without legacy applications to protect. In fact, in The Netherlands, with the Gigaport project, we saw this with the ADSL. There was no interest from the incumbents to provide us with decent ADSL solutions. So we now have two projects with two start-ups in different parts of the country. And they are bringing us quality ADSL for reasonable prices. And this forces the others to react.

COOK Report: But you do this by putting out tenders for a certain kind of ADSL and if the big companies don't respond to the tenders, that's their problem, not your problem.

Neggers: Yes. We got responses proposing inferior ADSL. We rejected them. They didn't even qualify our specifications, so it was a non-bid, you could say. And the same with this ten gig lambda network that we are building now. When we tendered that, it was absolutely impossible to get these kind

of offers from the large telcos. They offer you ATM and SDH. And we said we don't want that.

COOK Report: Because they've invested hundreds of millions of dollars in building up that infrastructure.

Neggers: That's not my concern.

COOK Report: Too bad if you're stuck with it, right?

Neggers: We made it mandatory that there would be nothing between the router and the lambda in the tender. So we didn't even ask for free solutions how to solve the problem. And then it turned out that we got five offers also from the incumbent providers and now BT is providing us with this network.

COOK Report: The BT subsidiary?

Neggers: Yes, Telfort, the BT subsidiary. And by doing this, we cause others to invest in laying fiber conduits to us. There is more invested in going to our POP locations than we pay for the network. Also industry is investing more than we in this network and the results are commonly available for everybody to use.

COOK Report: So the winners of the ten gigabit lambdas to your routers had to lay some new fiber.

Neggers: Yes.

COOK Report: And the whole point is the fact that you are laying new fiber from point A to point B, is publicly announced. Everyone can jump on your scheduled deployment of fiber and put their fiber in.

The Dutch Equivalent of Fiber Brokering

Neggers: That's what happens. There's an enormous investment in these fifteen locations compared to the rest of The Netherlands. I'm sure that this investment would not have been done without our tender and no one would have known as soon that industry was able and willing to provide routers and lambdas that could nicely work together at these speeds.

COOK Report: I see now why my friend doesn't like the American university CIO way of doing things. Because all they want is bandwidth and they don't care what kind of bandwidth they have and they just want somebody to subsidize it.

Neggers: But now you see also the advantage that it's not just the Ministry of Science and Education that supports the Gigaport project. It's also the Ministry of Transport, the Ministry of Economic Affairs, the Ministry of Municipalities, etc. Because we stimulate the deployment of conduits and dark fiber.

COOK Report: They have their own communication needs and they can take money that they would make their own tenders for use it in a way that is synergistic with what you are doing?

Neggers: Yes, because once these fibers are there, and certainly if there are competing fibers there, that's what's happening. Because once you have three or four or five competing operators, the prices become reasonably close to cost. And that's the enormous difference from the past, where you had no real competition. There was a single supplier and they dictated the terms. They decided when it becomes available and for what price.

COOK Report: What is your advice or prescription for what the intelligent user and the intelligent investor in these services should be looking? I mean, who do you think really has the answer and what are the problems, the pluses and minuses? What do you think should happen in the next year?

Neggers: My personal view is, the network comes first, the applications come later. Nobody starts working on advanced applications if they have no guarantee that a capable network will be there to carry it. Secondly, all networks will be based on fiber in the future. So the trick is to get in the dark fiber game as soon as possible.

I have gone through all these generations of networks, SURFnet1, -2, -3, -4. And every time, the usage of the next generation was completely different from what we anticipated. The users will find solutions for their problems.

COOK Report: Just give us the bits.

Neggers: Yes. MP3 was not invented by carriers. It was the users of the network, who started to do this. And it was a spin-off from earlier standardizations for digital video discs, etc. And it got a life of its own, because the network was there, as well. You remember the ASCII days of networks. Then you could have formatted documents. In the beginning people said there was no need for formatting. Just ASCII emails were good enough. Then formatted documents became common. The process was assisted by the emergence of a single standard.

COOK Report: MIME and that sort of thing.

Neggers: Yes, so standards are also important, but standards also mature, because sometimes there will be a critical mass. When this happens that's it. It's done, and the rest will disappear again. And so we actually progressed from ASCII to formatted documents to pictures to audio and now we're going to have video.

Well, MP3 was a spin-off from the audio-video industry, for audio encoding of video discs. Now you have MPEG4, which is spin-off of MPEG1, which was the video disc, and MPEG2 which is the DVD disc. MPEG4 is a smarter way to compress video. So with MPEG4 encoding, you can have a complete movie, one or two-hour movie, on less than a CDROM. So let's say 300-400 megabytes. One MP3 song is about 3 million bytes. So a one-and-a-half hour movie is a factor of hundred larger than a MP3 song. That's equivalent to the growth factor in bandwidth that we see in our move from SURFnet4 to SURFnet5. And that we used in the move from SURFnet3 to SURFnet4. This factor of hundred in 4 years is in fact a growth factor of 3.2 per year. This is more or less what we have seen in all our generations of networks. So with SURFnet5, it will be as easy for a student to move a MPEG4 encoded movie on our network than it is today on SURFnet4 with an MP3 song. And of course, this will happen. And, of course, industry will have to adapt to that. And you will be amazed at the quality of these movies. You can watch them full screen.

COOK Report: And as far as convergence goes between Internet telephony and the public switched network telephony, if you don't have to worry about Quality of Service issues because you've got lambdas to throw at the circuits, that opens that bar wide, right?

Neggers: Yes. The Quality of Service issues disappear completely. Doing video-conferencing at the moment on our network, is not a problem of bandwidth. The problem is you need hardware on your desktop or laptop to do the encoding and decoding. In particular the encoding. If you don't have that, you have lousy frame rates, small pictures. As soon as you have hardware inside your desktop — and you don't need more than $600 or $700 devices to do that — you have marvelous quality and the network is absolutely not the problem. Even without Quality of Service today. And what we want to do, of course, is prepare for the next step to be in time when all these people have

hardware encoding in their desktops.

COOK Report: That leapfrogging again.

Neggers: We don't want to become the bottleneck.

Wiring the Student Dorms

COOK Report: All the time then you are taking the government's investment and using it to leverage the building of new infrastructure. For example can you describe for me your cost model on wiring the student dormitories?

Neggers: Yes. One of the things we did to stimulate the introduction of high speed access for students from home is to make a detailed study what it would cost to wire a dormitory or a number of dormitories in a certain range, with gigabit Ethernet connections to our points of presence. And for the calculation we used a distance of 10 to 15 kilometers, switched Ethernet boxes in the buildings and Ethernet cabling with hundred megabit access to the rooms. It turned out that if you take all these costs and derive the depreciation for the fibers, for the switches, for the cabling, it turns out that it would be less than $20 per month per student, much less than that.

COOK Report: And for less than $20 per student, the resources you make available to the students are?

Neggers: They will have an Ethernet RJ45 socket in their room via which they can connect to the network.

COOK Report: And it's for on-campus resources, a hundred megabits per second? Do I understand that right?

Neggers: No, they will have full Internet connectivity from their rooms. Our point of presence is at the university as well, so they connect either directly to us or via the university network.

We made available a million guilders, which is about $400,000 U.S., to stimulate the development of these plans and investments. We got responses from several university cities. We granted the contract to the University of Delft, which is now wiring about 3,000 students in town. They have dark fiber from a small start-up to connect the buildings. Contracts with equipment providers to install the switches and to do the wiring they have another contractor. This will all be operated in a joint effort with the university, which is subsidizing the long distance stuff, and the building owners, who are just selling it as part of the normal rent of these rooms. And because of this little subsidy of the university for the gigabit Ethernet connections, the dark fibers, the cost for the Delft people will be even less than $10 a month. And, of course, they will have then full access to the SURFnet backbone and from SURFnet to our international connectivity. And as a result these students will have a different look of networking than others, because the network is a rich and always present commodity. They will start to work differently and to think differently and this is exactly what we want to stimulate.

Now let's move from an examination of the dormitory infrastructure to the general fiber infrastructure within our communities. If you want more than, let's say two megabits per second, your copper wire to the central office is no longer good enough. So the provider has to install a fiber to your premise to give you more than the two megabits. This is going on for years. Every provider who does this will charge you an initial fee which is related to the distance that must be trenched in order to bring the fiber to you. And then, in fact, you are locked in to that provider for the duration of your contracts. This is not a very economical way. Not for the provider and not for the consumer.

Therefore it's obvious you have to economize on laying dark fiber. The way to do that is to join forces with your neighbors. This is exactly what Canarie showed with these schools and municipality buildings. They contracted not with a telecom carrier, but with someone to dig the street, lay the fiber or put it on the poles. And own the fiber. And then start to look for a provider to connect to that network. The availability of the gigabit Ethernet standard and cheap interfaces makes it even more simple for anybody to run and operate a regional network and provide it to the upstream providers.

The fact that you can have VLANs on such a Level 2 infrastructure makes it even possible to split communities and connect them to different providers. For instance, the part that are allowed to take advantage of our subsidized network can connect to us and others can still connect to the commercial environment. You can even team up with commercial customers and your own government participants. And this is a very strong model. So the future is really in owning your own dark fiber, going to a place where you can meet different upstream providers and then you will have a competitive market and then you will have a market that copes with the available technology.

COOK Report: And having national governmental policies that encourage doing that where you're building new tracts of houses, and new communities.

The Public Policy of Fiber Infrastructure

Neggers: In the Netherlands, it's completely liberalized. Everybody in the Netherlands is allowed to dig fiber in the ground and own it. There are few restrictions. If you sell your services to the public at large, then the municipalities are not even allowed to ask you a price for it, they have to allow you to dig, because the only privilege they have is they are allowed to coordinate the digging.

Note that not everyone is allowed to dig every week. And, because of this, if one provider asks permission to lay a fiber, it will be announced and all others who are interested in the same route are free to use the same digging and then the providers have to share the costs of the digging together. So for the original one it becomes cheaper. For the other ones, they only pay part of the excavation.

COOK Report: So a fiber dig is scheduled thirty days in advance? Or sixty days in advance?

Neggers: Normally, municipalities don't allow you to dig more than twice a year the same street, or even less.

COOK Report: Makes sense. But everybody knows the rough period on the calendar when you can do this sort of thing and so you can coordinate and plan.

Neggers: And, of course, data communication is completely free, there's no regulation at all. As far as the government goes, you only heard part of my talk this morning. I talked about what is called the Digital Delta. This is the official government policy in The Netherlands. They want to have The Netherlands to be fully part of the information society as soon as possible. And the white papers that describe the policies are called the Digital Delta. And last month they gave an update, called the Digital Delta Beyond eEurope, where they noticed that, although we are seventh ranking on the ICT evaluation, extra effort may be needed to maintain a place in the top ten. One of the extra efforts they are doing is to stimulate the widespread availability of dark fiber. In particular they are also to start pilots for schools, hospitalities, municipalities, etc. They really are going to use this model as an official policy.

Light, IP and Gigabit Ethernet

A Road Map for Evaluation of Technology Choices Driving the Future Evolution of Telecommunications

Contrary to some opinions, the *COOK Report* finds that the Internet revolution is not spelled dot com. The revolution is in fact to be found in a total revamping of the transport of bits. While the dot com empires of 1999 collapsed in 2000 the cost effectiveness of pushing the Internet Protocol over glass yielded more dividends than ever before.

A growing amount of telecom traffic has migrated to a growing amount of fiber. The pure Internet play throws out SONET effectively doubling available fiber in the case where redundant loops were used. Whereas lighting each new fiber used to call for new bays of OC-48 SONET equipment at $100,000 a bay and up, a strand can now be lit at a gigabit by a $7,000 Ethernet switch on each end.

While gigabit Ethernet over glass is the current preferred Internet way, ten gigabit Ethernet transport will be arriving by year's end. If 40 lambdas per strand were high end in 2000, 160 is likely to be common by year's end. With the completion of multiple metro fiber build outs, end-to-end fiber may now be taken for granted by most business customers. The explosion of bandwidth as the result of more fiber and technology that squeezes more bandwidth from each strand has meant that, in some instances, the delivery of a gigabit costs about what a T-1 did a decade ago.

The bottom line is that telecommunications which is prepared to forego traversing the legacy PSTN is now upwards of 1000 times cheaper than that which powers a circuit switched voice call. While corporate managed VPNs have been able to avoid the PSTN for some time, a new development has emerged in Canada where customer management of optical wavelengths using the OBGP protocol holds the promise that by year's end users of Canada's new public sector gigabit Ethernet over fiber infrastructure will be avoiding carrier clouds entirely.

At the basic levels of both transport and network management the Internet revolution is shaping up to tell the PSTN that it is no longer needed. In telephony meanwhile protocols are being developed that will allow the diversion of large amounts of PSTN traffic to the Internet. ENUM is the major such protocol. This will allow Internet carriers to offer and deliver many services to PSTN attached phones that the PSTN itself cannot negotiate. Other protocols such as instant messaging are shaping up as coordinators for PSTN activity and off on switches that can control Internet connected devices.

Fiber to the home is becoming more common and companies like World Wide Packets are gearing up to make gigabit Ethernet termination equipment that will give connected families, telephone, fax, high end video, ordinary TV and data off of the same line. Canarie the Canadian advanced internet agency has some interesting ideas about these developments stating that Divergence rather than Convergence may be the key to low cost fiber to the home. Here is a narrative paraphrase of the language of a slide from the presentation 'Optical Communities' in September 2000.

When people first started looking at Fiber to the Home (FTTH), they deemed it to be too expensive because it assumed all services would be converged – date, voice and video. They noted that expensive terminal equipment would be required to segregate voice, data and video services at the home. Meanwhile voice traffic has largely gone wireless. Note that lifeline voice can significantly increase system costs by demanding high reliability and depending for this on DC battery power, 911 services. Perhaps it is time to conclude that the big driver for residential broadband is not voice or video. It is the Internet. Very soon Internet will carry video and second line voice. So instead of building a converged network such as FSAN, HFC, etc build an Internet network only. Divergence rather than Convergence may be the key to low cost FTTH.

While the power of the new systems is awesome, there are additional issues that will keep very interesting the life of anyone who must evaluate these changes and plan a winning strategy for the future. While one better be aware of the key differences in the power of the technology when compared to the circuit switched world's way of doing things, one also needs to understand that progress has, in this case, waded out into new and uncertain territory. There are some growth and scaling issues where the answers are not yet clearly understood.

For example readers should consider Bill St. Arnaud's paper on scaling issues of Internet growth. <http://www.canet3.net/library/papers/scaling.html> If the suppositions in this paper prove to be correct, then the role of backbones will have to be rethought and much Internet topology rearranged. One of the developments that influenced his thinking occurred on December 4, 2000 when Essex Corp announced an advance that enabled it to push as many as 4,000 lambdas down a single strand of fiber. On December 8 St. Arnaud posted the following to his Canarie mail list.

"[Here are some excerpts from a recent article in Lightreading that I believe illustrates the point that Ultra dense wave division multiplexing systems is not about bandwidth, but connectivity. A number of companies are starting to realize that today's Internet architecture (which is still fundamentally based on the old telco network model) will not scale. As intelligence moves to the edge, the existing network architecture must grow at some function related to the square of the number of users. However [the problem is what to do] if the increase in the number of the customer's own wavelengths and the fiber in the network only grows linearly? Intriguingly such a customer owned network architecture starts to closely resemble the neuron architecture of the brain. Perhaps mother nature long ago discovered the most efficient architecture for distributed intelligence? — BSA]"

From the December 4, 2000 issue of Lightreading.... http://www.lightreading.com/document.asp?doc_id=2760

"We don't need thousands of wavelengths for bandwidth, we need it for connectivity," says Simon Cao, chief technology officer at Avanex Corp. (Nasdaq: AVNX), a company that's developing high channel-count wavelength systems. Cao figures that the availability of thousands of channels in combination with tunable lasers will make it possible to eliminate complex optical switches, by using wavelength routing instead . That idea is the driving force behind all of Avanex's developments, he says."

"Despite its careless capacity talk, Essex has also thought out how to take advantage of the extra connectivity. Moodispaw reckons that Essex's solution will be perfect for customer-owned wavelengths in the access network. About 1 Gbit/s would be adequate to supply a gigabit Ethernet line to a business.

Each customer pays less for its wavelength, but the service provider is able to sell to a lot more customers, so everybody is happy."

What this means is that we likely soon see customer owned wavelengths of a gigabit. Such wavelengths will be plentiful, affordable and matched to the speeds of the Internet's basic gigabit Ethernet end-to-end architecture. When only the rich and mighty carriers owned fiber and the very expensive lasers that could pump bits and the even more expensive SONET equipment that sustained the bits, we imagined that bandwidth was a valuable commodity obtainable only from "on high," or from upstream. Indeed the carriers recognized their position and charged large sums for the scarce commodity that they delivered.

This situation has dramatically changed. If you own a piece of fiber in a metro area network, you can power that fiber with a lambda delivered as gigabit Ethernet for a one time investment of about $15,000. When you run out of bandwidth, adding the multiplexer necessary to turn the one lambda into four is also affordable. In this sense bandwidth can now be generated by customers and used very cheaply at the edge of the network over hops that in ideal conditions can be as much as 100 kilometers. Suddenly these customers can throw cheap bandwidth at quality of service issues. Their whole outlook on life starts to change rapidly. George Gilder was observed to say that "the companies that exploit bandwidth recklessly, will win."

On January 29, 2001 Bill St. Arnaud observed "Although I may disagree with Gilder on some of his predictions, I think he's gotten this one right. To date most advanced optical network research is based around the assumption that bandwidth is a precious and rare commodity and therefore we must optimize the network topology and try to extract every possible bit of capacity. One of the underlying assumptions of the existing CA*net 3 network and proposed CA*net 4 network architecture is that bandwidth can be wasted. Rather than concentrating on applications and technologies that optimize the use of bandwidth, we want to concentrate on applications and technologies that waste bandwidth but in doing so derive the maximum benefit for the user - hence our focus on "customer empowered" optical networks. There is no question in our mind with customer empowered networks that bandwidth will be wasted, wavelengths will be orphaned and seen from the perspective of a central carrier the network topology will be less than optimum. But in a world of thousands of wavelengths and near infinite bandwidth who cares? The true measure is whether the customer can easily and rapidly deploy new broadband applications and services."

"At CANARIE we are initiating a number of projects with this premise in mind - OBGP which will allow customers to setup and tear down their own wavelengths at will, WDD-wavelength disk drives will drive new concepts in distributed super computing, object oriented networking enables the wavelength to become a software object, community based condominium fiber networks, and much, much more. Stay tuned for more announcements and developments in this area."

Arnaud's words portend a seismic shift in the telecommunications landscape. With the exception of some important qualifiers we believe that he is correct. One must understand that, for the time being, bandwidth is cheap only in proportion to the distance that it must travel. Very high bandwidth sent over long distances is still quite expensive. Bandwidth that must be provided on a corporate wide scale for mission critical activities is also somewhat expensive. What Arnaud is talking about will change the world. The only question is how fast.

To sum up. Costs are falling. Depending on the regulatory environment, the amount of capital necessary to become a small scale telecommunications provider is rapidly shrinking. Independent back yard gigabit Ethernet providers can deliver broadband services at a fraction of the cost but equal in quality to what is accessible from larger more traditional companies.

There appears to be a choice of two paths to our telecom future. One is to go with the highly innovative pure Internet approach of gigabit Ethernet over condominium fiber. Such a choice empowers the customer, facilitates decentralization over centralized control and provides small and innovative businesses with the environment that they need in order to flourish. The other path is to try to fore stall the innovation by squashing competitors with a massive vertically integrated company founded on older technology and leveraging access to content and over a network monopoly so pervasive that people will find they have no choice but to buy it. What could be in store for us all, if things go in this direction, was summarized by Scott Cleland CEO of the Precursor Group on Friday January 19th, 2001. "Precursor believes AOL-TW has budding 'Microsoft-like' potential to grow increasingly dominant by being the leading national company that brings together the various online interfaces (content, Internet access, buddy lists, instant messaging, etc.) to become the de facto consumer online access market standard much like Microsoft Windows brought together the various desktop applications to become the de facto consumer software market standard."

On the one hand, the AOL infrastructure, which, compared to what the Canadians are doing, can most kindly be called archaic, has to be a terribly attractive vision for the ILECs. After all, it will give them plenty of time to finish depreciation of their copper plant. On the other hand innovative Internet technology ventures have the most to gain from taking the Canadian road. To facilitate making this choice with more certainty, in this publication we offer a recapitulation of *COOK Report* interviews on the shifting market domain of the optical network. These interviews are unique in that they treat complex key technical issues with attention to depth and detail found no where else. We believe that they form a set of resources that will permit readers to decide for themselves what direction is most beneficial to the future of their enterprise.

As we ponder these developments, we realize that we may be standing on the cusp of events that may overturn the economic structure of the first five years of the commercial Internet. The largest backbones, UUNET and Sprint, set the price of bandwidth and determined conditions for connection to the Internet. Last week WorldCom announced that it was debranding mighty UUNET and Mike O'Dell, the chief architect of the Internet's largest global backbone, left the company. The problem is that these services take more and more resources to pay for ever larger routers and support huge growth rates in traffic. It is beginning to look like revenue derived from running such a facility may no longer match expenses as the availability of alternative sources of cheap bandwidth and Internet transit increases. For the largest backbones sources of power only a year or two ago may now be turning into liabilities. The economics of bandwidth are in flux to such an extent that there is growing suspicion that for many large players business models that were sustainable five years ago have lead to the acquisition of huge debt today and that the next year may bring a round of major bankruptcies as the new economies of the technologies described in this report render traditional business models based on old technologies unsustainable.

In the end the path of the internet as the stupid network driven by these new technologies will predominate. We begin this six part report http://cookreport/lightipgige.shtml with a section on the basic Internet architecture and technologies of "Light, IP, and Gig E" to be followed by a section on their adoption in Sweden, Holland and of course our long special issue on Canada.

Bandwidth, User Tools Migrate Toward Network Edge Fueling Idea of Always on Disciplinary Computational Grids

User Control of Bandwidth Raises Interest in Shaping Network Environment to Needs of Subject Matter Communities

Editor's Note: Peter **MacKinnon** is Managing Director, Synergy Technology Management, headquartered in Ottawa, Canada. He is responsible for an international business practice with a focus along two fronts. The first is in support of organizations that are seeking to use Information Technology in strategic ways in order to transform their business. The second involves working on the interaction between Information and Communications Technologies and public policy development, with particular reference to the design, development and evaluation of programs intended to stimulate knowledge-based economy developments. Recent work in China has been to assist corporations, primarily in the Shanghai area, to better understand advanced Western management practices, particularly with regard to the Chief Information Officer (CIO). This ongoing engagement is part of a broader effort to assist Chinese companies and government organizations in preparing for their country's entry into the World Trade Organization (WTO). In Australia he has been assisting a number of organizations with plans for the development of advanced broadband networks at both regional and national levels. This is part of a national Australian initiative to provide the scientific and technical community throughout the country with state-of-the-art research networks, along the lines of the capabilities of Internet 2 in the U.S. and CA*net 3 in Canada.

His degrees are in earth sciences. He worked as a scientist for a number of years before becoming involved in international scientific affairs as the manager of a paleo-climate program in a U.N.-sponsored agency dealing with climate change. He has also served twice on assignment to the Government of Canada. On the first occasion he served as the Special Advisor for Advanced Information Technologies within a central agency and on the second occasion he was assigned to the Diplomatic Service in a posting to the United Kingdom where he was responsible for a group dealing with foreign investment, scientific affairs and strategic business alliances between Canada and the U.K. In between these various endeavors he headed corporate development for Cognos, a billion dollar software company, and was the founding president and CEO of the largest company in Canada to specialize in artificial intelligence systems. He founded Synergy Technology Management in 1994. We interviewed him in Montreal on November 29, 2000 on the occasion of Peter's role as chair of the Grid Session at the Annual Canarie Advanced Networking Workshop.

A post interview comment from the Editor. In talking with others about grids we have heard the following concerns expressed. A grid in a general sense is a seamless mesh into which anyone can plug to receive ubiquitous computing capabilities. However, various people have different ideas about implementing them. Internet 2 sees grids as a physical infrastructure limited to Internet 2 members. But we suggest that this is tantamount to an exclusive club like approach and is therefore contrary to the basic nature of the Internet. The second part of the grid concept, for many of the players at least, focuses on application layer changes that may be installed over either of these approaches at the networks basic transport infrastructure layer. Such changes can make a given environment congenial to physicists or genome researchers or some other discipline.

COOK Report: Grids are in the 'air' these days. What are they all about?

MacKinnon: A Grid is a metaphor for something that conjures up in the mind's eye some kind of mesh of interconnections and something flowing through the nodes of the mesh. In effect, the metaphor is in respect to the electrical power grid. The original electrification process, as you know, started with power generation and local power lines radiating from the generation station. Those who consumed electricity were nearby to the power plant because there wasn't any distribution system. Moreover, there were no links between different power generation networks. Once various power suppliers integrated their distribution systems you started to see a significant transformation in the economy, in terms of where you could be physically located, whether it was closer to your markets, or closer to your supply, depending on the economics involved. Even more importantly, the electrical power grid brought with it reliability in supply and standards in performance (i.e., standard voltage and current). This resulted in the rapid development and use of new technologies, products and services that depended on a reliable electricity supply.

COOK Report: All of a sudden, it was more ubiquitous and more reliable.

MacKinnon: Yes and when there were outages in one place, electricity still flowed, because the grid was interconnected. So the metaphor in a computer networking or communications sense, is that you have ubiquity and you have a high level of reliability. You also have the multiplier effect of having now established this capability, for with it, you can do many more things than you could possibly have done before. The distribution system itself adds value. The result of these new electrical 'marvels' transformed our economy and way of life. As far as grids go we are still in the early days. nevertheless I and many others believe that the promise of digital grids could be as transformative as its metaphorical predecessors.

Varieties of Grids

Beyond the metaphor and expressed in contemporary terms, a grid can be characterized as a scalable, wide area, broadband network linking a combination of data collecting and data handling capabilities with or without a computationally intensive resource pool and support infrastructure distributed across the network.

There are different kinds of grids. However, the differentiation is a matter of degree. Basically, the underlying infrastructure is comprised of networks and process devices (e.g., computers, scientific and technical instruments, and data storage and display devices) . For example, a 'data grid' is primarily concerned with gathering, storing and disseminating data, very large quantities of data. Particle physics is a classic field that could benefit from globally distributed data grids.

Typically, major scientific instruments are limited in number yet have globally distributed users. I could elaborate more on this theme latter.

You can have 'computational grids'. In this case the grid is viewed as a network of distributed computing resources that can work both cooperatively and independently of each other. They allow applications to operate in a distributed multi-platform environment across various geographic scales defined by the physical networks involved. Computational grids represent one of the frontiers of computing. They raise many fundamental challenges in computer science and communications engineering, much of which has to do with partitioning a problem across multiple machines, latency in the networks and administration and allocation of the grid's resources.

Then there are what I call 'access grids'. This is yet another way to look at grids. In this case a grid provides remote access to devices, such as, radio and optical telescopes, be they space-based or on terra ferma. Such a grid could allow a remote user to steer one or more devices to run and experiment and gather data (e.g., aiming, focusing and tracking an object with the Hubble Space Telecscope). This approach provides a dynamic real-time interaction between the user and the device.

Grids are relatively new and still rapidly evolving in terms of capabilities and applications. Today, most applications address scientific and engineering problems, hence the references from these fields. Furthermore, grid developers and grid users tend to be operating on the frontier of developments. To date, grids have not yet emerged in the commercial world, although several entrepreneurial groups in the U.S. are attempting to offer limited grid services.

Also, one can envisage applications in business and government. Much of it might relate to accessing information, particularly the way in which you could pose questions to mine information in the future. Furthermore, corporations might well use grids to develop products and support customers. Regardless, in the scientific and engineering worlds, we have this nascent and evolving environment we call grids. The building blocks for grids are also evolving. These include computer and communications hardware and software and the ways and means of combining these into serving practical 'business' models that make grids worth using. For example, in the communications area, everyone is of course moving to adopt fiber and the bandwidth that comes with it.

Let's look at broadband fiber networks for a moment. There are whole new business models emerging. When you build a fiber network, there are clearly significant economic returns expected. These new business models include such approaches as long-term dark fiber leasing and multi-partner condominium fiber builds. The grid frontier, going back to the science problems, is a dependency on broadband networks that then have various types of devices scattered throughout the networks, in the context of an interconnected grid.

A big challenge in grid development is partitioning and synchronizing distributed computational processes and data requirements. By example, suitable uses for grid applications fall between two classes of problems. One class is basically how to solve a computationally intensive problem distributed across a heterogeneous mix of computing resources in which very little data input is required but there's a tremendous amount of computing required (e.g., calculating the energy regime in a chemical reaction). The other class being problems that require highly intensive computing and a lot of data. For example, problems involving both lots of data and lots of computing include genomic and climate change studies.

COOK Report: In the scientific models, in the last ten years, you've gone from single supercomputers to groups of computers that are doing parallel computing of some fashion, right?

MacKinnon: Indeed. This is really where the grid fits in as a new paradigm for discovery. Let's explore some examples by starting with the SETI @ Home project. The search for extraterrestrial intelligence (SETI) is a monumental challenge. In the case of SETI@ Home the objective is to use an Internet site (www.seti@home) as a source for downloading both the computing software and associated data sets via a screen saver. Typically a screen saver is automatically activated when local computing requirements are idle and some type of graphical display is used to indicate that a machine is on but not in use. In the case of SETI @ Home, once the screen saver is downloaded, each time the machine kicks into idle mode the screen saver software begins to search a package of data extracted from the project web site. Once data have been processed, results are automatically transmitted back to the project server and a new batch of data is then automatically downloaded for the next round of processing. This rather simple procedure has lead to one of the most remarkable experiments ever conducted on the Internet. In less than two years hundreds of thousands of individuals all over the world have signed up to process SETI data. The result to date is that over 186,000 CPU years of computation have been devoted to searching for SETI signals and most of this has been done by volunteers. So far, ET remains elusive.

As a sign of things to come, SETI @ Home operates like a grid of peer-to-peer machines. A commercial extension to this notion of peer-to-peer computing comes from products like Gnutella and Napster, which have typically been used for music downloads. However, when considering the underlying functionality of these types of products, one can see a paradigm shift developing in the operation of the Internet. The change is a dicet linking of machines to each other whereby any number of machines may share processes and memory directly with each other. Who knows what killer apps might lurk in this emerging peer-to-peer world.

Bringing Together New Resources and New User Populations

Grids offer a new way to optimize the use of scare resources while allowing large numbers of researchers news ways of sharing and discovering. Thus, SETI @ Home is an example of a whole new way of conducting scientific research. This has been brought about by the grid-like function of the data analysis component of the project coupled with the ability to allow vast numbers of individuals to work on and share a common scientific problem.

We have some related parallels emerging in the commercial world. For example, Cisco, which has been extremely successful as an Internet equipment supplier, has evolved one of the most effective e-commerce web sites at <www.cisco.com>. Basically, their site is customer-driven, especially when it comes time for customers to specify their needs. As I understand it, Cisco allows its customer to specify and configure their needs on-line through an elaborate business portal. The point here, at least in terms of grid developments is that Cisco uses their website as a concentrator of ideas and needs as expressed by on-line users worldwide. Although the actual application is not a grid in the sense we have been speaking about in the scientific and technical community but it does have a likeness to a conceptual grid in so far as customers or users scattered across the globe access data remotely, interact with the data at the remote location (i.e., Cisco's web site) and download results to their local machines. At the same time, the accumulated

impact of customer stated requirements and configurations and the like allow Cisco to capitalize on new ideas and common trends as defined by the 'volunteers' who use the site. In fact, there is a loose analogy here with the SETI @ Home project.

COOK Report: I decided to check out your assertion and at http://cisco.com/warp/public/779/ibs/ found " Welcome to Cisco's Internet Business Solutions. Cisco has set the standard for business transformation by using Internet technology to integrate its core processes and culture. The result has been dramatic: monthly online sales top $1 billion, annual operating cost savings are about $1.4 billion, 70 percent of support calls are now resolved over the Internet and customer satisfaction has increased significantly."

MacKinnon: A more complex example called "Neptune a Fiber Optic Cable to Inner Space" is a proposed grid that would be deployed over the surface of a tectonic plate on the ocean floor. This initiative is being led by Woods Hole Oceanographic Institution and NASA's Jet Propulsion Laboratory (JPL) in the U.S. and the Department of Fisheries and Oceans in Canada. The proposed project is currently advancing through feasibility studies. Basically the intent is to put a large fiber optic grid on the Juan de Fuca plate off the West Coast of North America. The grid will have many nodes spaced tens of kilometers apart . Each node will have a standard instrument receptor that will allow for a wide range of monitoring equipment to be linked to the grid. The site has been selected as an ideal location to conduct a wide range of oceanographic and geophysical experiments and environmental monitoring.

There will be a group of providers who will put the infrastructure in place, and then the scientific community can design, develop and plug in their packages to specified areas within the grid that are of interest. In operational terms this project is both a data and access grid. Computational horsepower is left to terrestrial sites that can link to the ocean floor grid resources.

There is however at least one major problem on the horizon for this project remaining to be resolved. The Juan de Fuca plate also happens to be off the coast of Washington State where the U.S. Navy maintains one of its Trident submarine bases. Having sophisticated scientific instrumentation in such an area could, in the view of some, be deemed to infringe on national security. Nonetheless, this is a serious proposal which is still subject to more detailed study, including developing the engineering requirements.

COOK Report: By the same token, has anybody been thinking about putting that kind of a grid around the San Andreas fault?

MacKinnon: In some respects, this has been done in so far as deploying vast arrays of instruments that are sampled remotely -- a kind of simple data grid, if you like. In effect this is yet another example of how we are starting to integrate the data gathering with processing capability. As this integration continues and new functionality is added it will be possible to steer an application in real time with real data.

COOK Report: Give me an example.

MacKinnon: An example I can think of, is an application used for training in the area of computational fluid dynamics. It is called CFDnet and can be found at <www.cfdnet.com>. Computational fluid dynamics is concerned with examining fluid flow over or through some kind of geometric structure. You may have a very simple flow, like laminar flow. However, should you introduce an edge and then continue the flow process, you start to get eddies and so forth. The data and computational requirements for this kind of application grow exponentially as the complexity of the flow patterns increase and the resolution of irregularities increases.

CFDnet is a software package that has been engineered so that it uses a browser at the client end. The server is one or more high performance computers that perform all the intensive computations. This is done in an architecture that is similar to a computational grid. Depending on what your interests are, you can, in effect, change the dynamic of what you're simulating in terms of scale, almost in real time and depending on the amount of processing power you have available on the server side. You can actually steer your way through a changing two-dimensional model. There are also plans here in Canada to make CFDnet a research tool for the C3.ca Associates consortium <www.c3.ca>. This is a national non-governmental initiative to create computational grid facilities and capabilities in Canada. Plans call for expanding CFDnet to support 3D spatial capability and to then use it as a testing environment for various kinds of grid services being planned.

COOK Report: In the business area, is there any possibility of doing something like this with some kind of financial models?

Financial Grid?

MacKinnon: I'm quite sure that there would be ways. For example, I could imagine using distributed neural network-based systems to do predictive modeling with specific applications to financial markets. There are already companies that maintain systems like this to monitor markets and make forecasts in order to help decision-makers work their way through some of the complex transactions and financial tools that are prevalent in the market today. Typically, these are highly classified corporate systems. However, I would argue that security considerations for this type of grid application are beyond current grid development efforts. I can imagine at some point that these systems, which are self-learning, will continually add to their knowledge pool with human intervention as well as self-discovery. The implications of such capability are mind boggling both within financial markets and beyond.

You could interconnect a number of these systems within, say, one network. With such a grid, a global trading house, for example, could end up having its neural network system in London connected with that in New York, Tokyo, and Sydney. So now you're monitoring on multiple levels in multiple markets. If you have the latency problem solved, which may not be as big an issue in data passing networks as process passing networks, then you can run in real time while doing the necessary computations and communications between components of the grid. This approach could lead to the development of new types of financial instruments or new ways of hedging or new risk reduction capabilities. I suspect that it can be seen in a variety of ways. The financial area would probably be one of the first major commercial uses of grid-like capability.

Another example of grid-like capability could well be in utilities. These are organizations that have distributed systems already that want to use the grid, in a simple sense, to do status checks, self-healing, monitoring, or whatever the case may be. In fact, many utilities are interested in offering broadband network capabilities via their own rights-of-way.

COOK Report: If I'm a power company or I have money to invest in building power plants, can I use this to look at patterns of power consumption and figure out where a good place is to build new plants?

MacKinnon: I suspect that conventional practice of where to build plants is well-developed and so I doubt if this is going to help choose where to locate a plant. But I can imagine that, if there's an ability to track, either in real-time or some sample of time that's not too grainy, the consumption of

power by user within a region, then there may be some very useful research opportunities in customer relationship management that could be derived from monitoring such patterns of consumption. Although such an application is not particularly computationally intensive, at least in a high performance computer context, there would likely be significant data gathering, storage and analysis requirements.

COOK Report: It could help you figure your pricing.

MacKinnon: Well, it could help to develop different pricing schemes, so that you could offer flexible pricing (e.g., dynamic real-time pricing). There have been schemes like that in the past, where there's been daytime pricing and nighttime pricing that's been common for certain types of industrial consumers. This one would allow monitoring and optimizing pricing for all customers. This is very simplistic in terms of computational requirements, compared to what the scientific community is trying to address. It is more on the scale of the SETI@ Home 'grid', I suppose, in terms of computational and storage requirements.

What Do Grids Need?

COOK Report: Giving me an overview sense, where is all this going? What needs to happen and what do grids imply for infrastructure beyond, presumably, another customer for bandwidth?

MacKinnon: In this case the explorers are an ad hoc group, involving Canarie, the National Research Council of Canada and the C3.ca consortium, which is the high performance computing consortium of Canada. We are looking at how grids could change the way in which research is undertaken, at a national level, since we're looking at it from a Canadian perspective.

As we are looking at what needs needs to be done, we are trying to build a conceptual model that allows us to think more clearly about the relationship of these complex tasks to each other.

Our model takes the form of a chart. It's a very simple Cartesian system, of which on the conventional X axis, we have research couplings. Originally, we had facilities sharing. Then we had wide area distributed computing,of massively shared data, so you're moving out in complexity on how you can deal with the actual tools to use for conducting the research, in terms of computation, data and the problem itself.

On the Y axis, we have the way in which you can be organized. So you have the individual researcher or an academic department, or an institute conducting a particular kind of research project. Then we've got on another scale something like National Centers of Excellence. The goal is to have distributed requirements working on the leading edge of whatever the problem area and research area might be.

Then at a higher level, what we are doing needs to permit moving into areas such as Grand Challenges of the type that have been proposed on a number of occasions, in the U.S., for instance. The Grand Challenge notion was championed to argue for high performance computing and to argue for enhanced bandwidth and so forth, through both NSF and the military.

So you have formal organizational arrangements that are going up in what would be the Z axis. Along the Y axis, you have new types of technological paradigms. Originally, we had email and ftp. Then the web. Then we have Napster, and Gnutilla-type of peer-to-peer relationships. We have new kinds of things happening this way.

What we're suggesting is that if you think of this as a space, it is, in fact, what we'd call "grid space." And we're thinking, well, how can you conduct research in this space? What does this offer you as a new tool to undertake research at any one of these scales? From simple, individual researcher to a Grand Challenge-type problem. What kind of resources do you know? Where can those resources be? Of course, they could be almost anywhere, because you're now getting to the point of having a ubiquitous, reliable, broadband communications infrastructure. You do need special equipment at certain places in that infrastructure. It certainly doesn't matter really where it is, either, other than being defined by human requirements, typical concentration of people with skills to both utilize, maintain and conduct things with that equipment.

This chart is a conceptual piece, where you're looking at ways of trying to define the network for the community. Now, this is preliminary thinking, but it's a way of trying to define, in our case, thinking about what the next direction will be for Canarie. Canarie has been a very successful organization in terms of how it has both delivered and what it has delivered to its community of interest.

COOK Report: The success of these grids will depend on leveraging everyone's infrastructure.

MacKinnon: Indeed it will and it is analogous to the same motivating factors that drove the development of the original electrical power grids.

COOK Report: To what extent are there discussions on the best ways of doing that?

MacKinnon: In reality,the community is just starting to get into solving many of the technical issues of integrating current communications and computing advances with the architectural needs of grids. However we have to find ways of solving certain problems in both areas before we are going to be able to make grids, in what we might call a promise, the success that they might appear to be.

Technical Issues

For delivering the promise of grids, as demonstrated by the notion of grid space, which we just talked about, the following are likely. Grids will provide powerful, interactive, dynamic and flexible environments allowing for opportunities to create new discoveries, on a level of Grand Challenges. They will also allow for more R&D without increasing other resources as well as widen access and enhance educational uses.

So much for what could be major benefits from the promise of grids. But to actually implement robust grids is going to require a great deal more advancement in software systems than is currently available. So when you start to think about what it is that you're going to do in a grid-like problem-solving environment, then there are some really fundamental technical issues that need to be addressed.

For instance, if you're dealing with a computationally intensive problem, there have been a lot of advances made in development of parallel computing approaches in the last several years. These advances allow problems to be partitioned so that multiple parts of the problem can be simultaneously computed on different processors. Therefore, you have to understand in some detail the kind of problem that you're dealing with in order to know how to partition the processing. Let's assume that you have a problem of this sort and let us for the moment assume that it doesn't require much data, so it self-generates by calculation.

Recently there was a very good example. It involves an operations research problem that was first posed in 1968 and finally solved mid 1999. The so-called 'nug30' quadratic assignment problem took approximately seven elapsed days and an average of 650

computers on at least two continents to simultaneously work on nug30 in order to reach a computational answer. After 96 thousand hours of CPU time the problem yielded an exact solution. Now that the solution is known it is possible to incorporate the nug30 results into desktop applications.

COOK Report: For somebody who knows nothing about operations research, what would be an example?

MacKinnon: An example of an OR problem that can make use of the solution to the nug30 challenge is how to configure machinery in a manufacturing facility, for optimal throughput.

There may be many different parameters at play, but you can look at them from at least two manufacturing points of view. First, in the sense of a continuous process -- paper making, for example, or steel making or something where you've got a constant supply coming in, and second, as discrete manufacturing -- such as automobiles.

At issue is how do you optimally lay out machinery on the shop floor, where best to locate supplies and intermediates, and where the people fit in. This is a fairly complex problem. Once the nug30 solution was found it improved our understanding of how to deal with this kind of planning.

Indeed, deriving the solution to the nug30 problem was a kind of watershed. It represented one of the first successful efforts to combine widely distributed high performance computing resources in a computational grid environment for a sustained period, and it worked. Prior to this HPC resources had been used to achieve similar ends but generally limited to machine room scale distances. Certainly other applications have been run on distributed systems but rarely on the scale and within the architectural framework that was used to solve the nug30 problem. Now with vast stretches of geography interconnecting machines that are sharing processes and data synchronization is a critical issue to the point that the speed of light and electrons through the networks and processor matters.

There are many other kinds of problems and applications being applied to grids. For example, grids can be used to support distributed virtual environments. For instance it could allow multiple users at distant locations from each other to interact with the virtual environment in real-time (e.g., multiple users manipulating the same virtual object). This can be extended beyond visual queues to include haptics – simulating the force fields associated with touch, as in grasping a tool. The manifestations of virtual distributed haptic environments (e.g., surgeon practicing complex procedures in virtual space) are telepresence application (e.g. actually conducting remote surgery). Applying these approaches to widely distributed sites is an extremely complex process, especially from the network latency perspective. Network routing delays, beyond some very tightly constrained limits, with some element of the 'operation' – virtual or otherwise, could be 'life threatening' and therefore unacceptable.

Lets look more closely at this example. The haptics researcher's name is Kevin Smith. He's with the Commonwealth Scientific Industrial Research Organization (CSIRO) in Australia and is Executive Director of the Interactive Virtual Environment Consortium in Perth. As the name implies his lab is focused on creating desktop level distributed virtual environments integrating haptics and applied to medical and geophysical applications. For example, one of the projects under dvelopment involves simuulating minor surgical procedures on a human hand. This kind of work requires complex real time simulation modeling of human tissue such that simulated hand responds in a manner similar to what would happen in real life. Because his team's applications are being used in a distributed environment, issues of computational peformance, network latency and quality of service are critical areas of research with respect to the application domain.

COOK Report: The system might have gloves and you can stick your hands in and your fingers do something that's sent to a robot that can be a long way away.

MacKinnon: Yes, that's an extension of this in the sense of telepresence, that allows you can do robotic surgery.

But what I wanted to say is, if you take a scalpel, for instance, and you start cutting the skin, you feel it, just like you were really doing it. In order to do this the visio-haptics system calculates in real time the changing topology of the skin as it's cut. And it calculates the changing tension and muscle strengths and all those characteristics that are part of the actual human anatomy.

COOK Report: So you have the same kind of feedback that you would as if you were working on a single cadaver.

MacKinnon: Absolutely. Moreover, if you hit a bone in this virtual space, you feel it. The system is designed so you can actually have real feedback. It's all there. It's incredibly sophisticated.

But now you can distribute that. You can put it on workstations across town, in a metropolitan area network. Still, speed of light isn't going to be a big issue there. So you can actually use it as a training or as an advisory, you can imagine different ways that doctors might use this thing.

But if you then spread it across a greater distance, say 7,000 kilometers, now you run into latency issues, because you're constantly computing to be sure that the remote activities you're doing at one end parallel what is being done at the other end. In fact, two or more people can be working on the same body part or physical device, if you're modeling clay, for instance, simultaneously, with their own discrete pieces of that model. This is all part of a distributed virtual environment. And that's another kind of computationally intensive modeling in a virtual environment.

You have, of course, data-type problems where you get into even more complex issues. You may have computationally intensive issues, coupled with the question of the location of the data. At the same time you must know how much data you need, what files do you get the data from? What is the data model?

Solving Basic Problems in Resource Allocation via Middleware

Are the data, in fact, located in a number of data stores in different parts of the world? If so you will run into another range of problems, with respect to accessing data, processing data and ensuring that it passes appropriately. It doesn't take long to realize that underneath all of what we've been talking about are other kinds of problems that have been solved at the computer room level. These fall into categories such as resource allocation, and billing, authority to actually use the process. These are not very glamorous things, but they're of absolutely fundamental importance to allowing a grid to function, because even in the technical world, the scientific and engineering world, people are starting to pay for the consumption of the resources. Even though it's not a commercial application, many of the things that are going to be needed in the commercial world are actually going to have to be developed, simply to operate these things.

So one of the classic bits of terminology, then, is middleware. Middleware sits inbetween the application that you wish to actually run, like the visual haptic application, for instance, and the underlying infra-

structure itself.

Middleware deals with authorities for billing, security, and scheduling. It also actually deals with the partitioning and reassembling. So that's all part of this middleware of which there are various types of approaches that are being developed. <www.research.ibm.com/Middleware2000/> Probably the most widely known approach, although not the oldest approach, is known as Globus. <www.globus.org> Globus is a large international, although dominantly American effort to develop middleware for grids. And there's a variety of outfits I just talked about that are projects developing the middleware, from scheduling through to partitioning. There are a range of these type of projects.

There are other, alternative approaches. There's one known as Condor, which is basically an operating environment for distributed high through put computing. Rather than sitting on top of an operating system, condor itself acts as a grid-operating system. It is being optimized for distributed computing. <http://www.cs.wisc.edu/condor/ >

Another one that is also very interesting is the Grid Forum which is basically a talk shop for those who are trying to build grids aswell as those who actually want to use grids. The grid foum provides a common meeting ground from which to discuss the entire dimension of grids. They have organized themselves into several working groups. See http://www.gridforum.org/

This initially started in the U.S., then expanded with some international participation. Then the Europeans created an E-grid. For European Grid. Which is a Grid Forum, in effect. <http://www.egrid.org/str_glowna.html > What's happened is that the various people in different geographies — Asia, Europe and North America — have now come together and created one new organization. This is the Grid Forum that you can get access to. Charlie Catlett is currently the chair. Grid Forum started at the Supercomputer meeting in '97. So we're talking about basically things that have come into being in the last three to five years. I mean, obviously, some issues have been around. Condor has been around since I think the mid-80s, even early 80s. You'll always find early explorers.

COOK Report: But is it the fiber infrastructure and gigabit Ethernet as enabling tools that are bringing this to a critical mass where it becomes worthwhile for someone like you to tell someone like me about it?

MacKinnon: Yes, in part, there's a variety of infrastructure changes that are taking place that now make it practical to start looking at these kind of problems. We can do all this in a "gedanken" sense and not even worry about whether there's a real infrastructure out there. But now that there is emerging a real infrastructure, and the potential ways of using that infrastructure, as, for example, grid space, challenge all kinds of new ways of thinking about how you can both probe information itself, data. Ultimately, I suspect we can see that grids may lead to knowledge management practices. I mean, you're well aware of the development of the CIO in many organizations and now we're starting to see the chief knowledge officer emerge in certain kinds of organizations.

Emergence of the CIO

COOK Report: Like what?

MacKinnon: Well, like financial institutions which are organizations that are highly dependent upon information. Ask whether the product of an organization is information in which case they are more and more likely to have a CIO.

Nonetheless, knowledge management is important in all organizations, no matter whether they are government or belong to the private sector. I think that the early adopters in knowledge management are going to be found primarily in large organizations concerned with information assets which from their point of view they may consider their principal actual "property". This doesn't mean they actually own it in the intellectual property sense, but it's information that they manipulate to make money.

Middleware is providing a way of leveraging the infrastructure to do things differently than the way that they'd been done before or to do things you couldn't do before. I say this, because if you wanted to compute the answer to this problem, it would have taken 700 years. I mean, you just haven't got that much time, so you don't try to solve that problem. However, if you can solve that problem in 700 seconds, then it worth waiting for. But you need a new way of doing it. And what's happening is a kind of democratization of the infrastructure.

Coming back to grid infratstructure for the moment there are many kinds of expensive scientific instruments serving different communities of interest. These tools are spread around the globe and into outer space. These are devices, whether they be high performance computers or data-generating tools like accelerators for nuclear physics. The fact that you can gain access to the mesh of computing resources further democratizes how these tools may be used and perhaps even more importantly how data are integrated and analyzed. From the point of the governments — who in most cases pay for these things, — it starts to spread the opportunity for discovery among a broader community. And it also allows, through the infrastructure itself, new ways of looking at how to pose problems or pose solutions to the problems and therefore maybe get some further value added.

COOK Report: Are there protocols that are needed to further enable certain things, as part of middleware? Are there any IETF type things that are falling out of this?

MacKinnon: I'm actually not aware of what the Internet Engineering Task Force is doing in this regard, if anything at all. It may actually be a little early to try and look at creating a standard in that sense. But the Globus program is probably the most advanced in terms of general acceptance around the world — and it's not a large number of people we're talking about here, but numbers into the hundreds. Remember however it only takes a few people to make something happen for as long as you focus on going in that direction, you can go to the moon.

COOK Report: What are some of the events, people, players, things that you see happening in the next six months or a year in this area?

MacKinnon: Well, the Grid Forums are an event. There's a variety of workshops that are held on both sides of the Atlantic that have working groups that can be followed, passively or actively, by participating. Depending on one's level of interest, there's a lot of opportunity to follow (as both users and builders) what's going on if one has an interest in grids and that level of technical conversation.

Infrastructure Building to Support Grids

From the point of view of infrastructure building, certainly North America is further ahead than the rest of the world. We see already that Canarie, which is one of the leading networking organizations from the point of view of R&D and supplying R&D backbones, is definitely considering a grid-like infrastructure as its next generation network. And that has implications to selling an argument to the national government, on the grounds that it's going to have a socioeco-

nomic impact on the country. And therefore if it's going to have socioeconomic impact on a country like Canada, given that Canada is in a league of countries like the United States and other members of G7 or G8, certainly G7, then it's likely to have impacts on advanced economies in general.

COOK Report: And I think I saw a diagram of the proposed layout and you have almost a two-dimensional backbone. You have a backbone that has X and Y axes to it, right?

MacKinnon: Yes.

COOK Report: From a design philosophy, I guess a grid would be perfectly amenable to routing and routing-type decisions rather than permanent virtual circuits.

MacKinnon: In effect, we're really starting to pipe dream a little bit here. But we have a notion, for instance, with the ubiquitous availability of bandwidth. And, as we're seeing more and more fibers per cable and increasing capability in multiplexing data on a particular fiber, we're in a position to start thinking about if we had dedicated wavelengths, whether we could actually configure a grid-like application for some temporary period of time in a dedicated wavelength context.

COOK Report: This is the optical border gateway protocol stuff, right?

MacKinnon: Yes. So this is where your protocol notion would come in, there's a need to handle that kind of multiplexing. But there's no new need to change the infrastructure. Because the infrastructure is fiber. And that's going to be stable and reliable for years to come.

COOK Report: This shows yet another example of the insatiable demands on the use of computing and telecommunications technology. And it's presumably important to the providers of any set of infrastructures as evidence that demand is not going to run out any time soon.

Is there any thought as to what this implies for the economics or business models for infrastructure or network construction?

MacKinnon: I have to suggest that it is probably limited in the context of your question. Limited though in terms of looking at broad economic models and their implications at the level of the firm or industry. But some considerations are being made, arguments are being developed. For instance, in selling a grid like concept to the Canadian government as the next generation internet. There we're looking at some of the major areas where governments make large expenditures, such as health, education. And that's universal. Around the world, it's governments who generally pay the largest share of health and education expenditures

But what we can see is that with the advent of tele-health or tele-medicine, where you're starting to see digital mammograms, digital imagery, high-density digital imagery, you're able to start to create slices of that digital imagery. Each of those slices are starting to be measured in many hundreds of megabytes. As a result, you're going to be into very, very large amounts of data. Petabit-data sets, in terms of all the X-rays and all of the ultrasounds, etc., etc. Yet at the same time, if this is all digital information, then it can start to be used in ways that weren't possible or even imaginable before.

COOK Report: And some researchers investigating other treatment protocols could ask for the ability to prepare radiographic sets of lung studies.

MacKinnon: That's exactly where I'm coming from. I've been talking a colleague at a major British University. This person of the view that, within ten years, you won't need pathologists as we known them. The reason is that pattern matching will be so well developed in digital pathology that knowledge-based systems will handle the vast majority of diagnoses. You'll need specialists to handle the rather rare things, but you won't need the numbers of pathologists as we have it today. And that's going to be a fundamental change in less than ten years. Now, he may be well ahead of his time, but if you're setting out to be a pathologist, ten years isn't very far into your career, especially if you figure it's going to take you ten years before you actually start practicing. We're already starting to see thinking in professions as diverse as medicine that some of the things that these technologies are potentially able to do are going to have a profound impact on their profession.

So that's starting to happen, but it's not a global thing yet, it's individuals in different places that are advocating these. It is through events like this Canarie meeting that we commonly share new insights.

COOK Report: When does, a global bank, for example, or investment brokerage, such as Chase, Citibank, Merrill Lynch, some of the European banks they have to start thinking about the financial modeling capabilities of their telecommunications and computer systems. They are probably going to want to treat much of this as proprietary data on private networks. How does the grid concept fit into there? Can it be cross-company as opposed to industry?

MacKinnon: Well, I think it's going to be like a utility. I mean that grids will be somewhat like utilities.

COOK Report: In other words, a Citibank doesn't have to worry about using part of its infrastructure for this?

MacKinnon: No, they don't need their own global infrastructure, they will use the public infrastructure. But the public infrastructure will have security and reliability that is of industrial strength and will allow businesses to make business decisions without having to worry about the reliability of that infrastructure. Or even the cost of that infrastructure. It's going to be at a cost that allows businesses to go ahead and make business decisions without having to worry about the cost. While we're some years away from that, there are also large demonstration projects being assembled in grid context.

NASA, for example, has developed the Information Power Grid. Basically what they're doing is they're linking a number of their major centers, major labs, where they have high performance computers, into a grid-like environment. And they're investing a reasonable budget to learn about how the processes ought to work. There's all the software and all the advances in the technology itself, both in the communications and the processing that need to take place and it's obvious many spaces have to be filled in. We don't have all that middleware yet. And obviously, we'll probably want to have several choices of middleware ultimately in the marketplace. But at the moment, there are several experiments that have been underway, with reasonably funded and dedicated resources to try and build the pieces.

NASA will also work with those pieces, but they'll also try to understand some of the human dynamics within their own organization. Because this is a whole new way of thinking. They may well be trying to apply a grid space problem-solving environment to their own research problems.

COOK Report: I was thinking about interdisciplinary, in the sense of merging wireless and fiber optics and that type of thing. Maybe this is another one of the three axes on your chart?

MacKinnon: Yes. Wireless was on there. But wireless is just a conduit. Unfortunately it has limitations and its limitations are its bandwidth. I mean, wireless, no matter how efficient and effective it can be for a lot of

things, if you're bandwidth bound, you may just not get the bandwidth you need out of wireless. Nonetheless, all the data and instructions that pass back and forth between earth and space based instuments contitutes wireless traffic. We got large quantities of data coming back. Sometimes you have to stream it, but nonetheless, it comes back.

COOK Report: What about the concept of using the network to access data remotely?

MacKinnon: Absolutely. That's one of the key things of the data kind of grid, to be accessing data remotely. And that data would be, that's why I said you have data storage and data management facilities within the grid. We're talking about wanting to manage — I mean, you're not going to bother if it's just a few megabytes of information. You're talking about huge quantities of data. And there are new systems that are going to be generating petabits of data, just simply on an experimental basis, like scientific experiments. And, as I say, if you think about the digital logs of hospitals in the near future, they'll be requiring huge quantities of data storage and data manipulation capabilities.

COOK Report: Is it folk like operations research people who are going to be out there to advise hospitals how to organize their radiologic data in such a way that it can be put on networks and, at the same time? I mean, do you serve as an entity that talks to them and gets their requirements and then goes and talks to the storage people and the network people and helps put all this together?

MacKinnon: Within the medical profession, there are some new kinds of positions emerging. These focus on the information medical specialist. I think that's the kind of function that these people would probably perform, among other things that are happening. It is a fact that we're seeing more and more digital output from medical devices and more and more, therefore, digital records being created.

COOK Report: Yes and you have to somehow organize these into libraries and figuring out what you want to do from a privacy point of view.

MacKinnon: Indeed working with any data of a pesonal nature leads to a need to deal with all the classic issues of privacy and security of content.

COOK Report: And then organizing the results in a number of ways that other people can benefit.

MacKinnon: There's a lot of research and policy development going on in that area already. Independent of grids they're just looking at large data storage, the privacy issues, access, fee payment schedules and all these sort of things that are related to the tele-medicine, if you like, and delivery.

COOK Report: What are the issues involved for a subject matter field constituency who wants to have a grid or wants access to a grid? What are the issues they face in putting together the economic network resources necessary for getting the infrastructure that they want. I mean, is this part of the picture that you're dealing with?

MacKinnon: Well, it's part of the picture, but at this stage, it's primarily among the academic community. Obviously, vendors are interested, and they recognize that this is an emerging opportunity. They're obviously contributing in kind in many cases, because it's partly where their future lies. Be it as a disruptive technology or as, in fact, something that is very much in the mainstream, or, at least, could become the mainstream.

The infrastructure is primarily publicly funded so, or at least access to the infrastructure, whether it's a US West fiber cable or an AT&T cable or a Bell Canada, Bell Nexxia cable, doesn't matter. Someone pays to get access. The academic community gets preferential pricing. For example, there's no user fee for Canarie's backbone. That wouldn't stand in the commercial world, because where are you going to make revenue? It's because this is an enabling thing, and so it's underwritten by government. For the reasons that are the right kind of reasons government underwrite investments in determining or seeking ways to where the future is going.

In many respects, I characterize Canarie's mandate as to go out into the future and report back. And one of the areas of going out into the future and looking around and trying to report back is on grids. One of the reasons we're doing that is because for reasons defined by the conditions under which Canarie is unfolding, in the end of March 2002, the network as we know it will cease to have funding. There are other options. More funding to do as we're doing. Stop the funding and let people find other alternatives. Or replace it. Given the way in which technology is going and the growth in which Canarie has been making progress, it seems apropos to at least consider a replacement option. This is being done.

COOK Report: Well, I think that probably is a pretty good introduction. Many Thanks.

CANARIE To Build World's Largest Disk Drive

Editor's Note: This is likely the first press release that we have ever published. (It comes from the Canarie mail list.) It is certainly a mind bending idea that leaves us with many questions. like OBGP, it looks to be another extremely creative pushing of the envelope on the part of our Northern neighbors.

OTTAWA, February 7, 2001 -- CANARIE, Canada's advanced Internet development organization, today announced a plan to construct the world's largest "disk drive" called a Wavelength Disk Drive which will be constructed around wavelengths of light on CANARIE's national optical research network, CA*net 3. The wavelength disk drive will be more than 8000 km in diameter.

"For several years, researchers have recognized that harnessing the computing power of thousands of personal computers connected to the Internet would provide more computing power than even the largest super computers," said Andrew Bjerring, President and CEO of CANARIE. "This innovative project is intended to address one of the challenges inherent in realizing this dream: the difficulty of sharing large amounts of data efficiently among thousands of computers, each trying to communicate with the others."

Instead of having the computers send and receive data from each other,which slows the collaborative effort down to the slowest computer or slowest network link, the computers will simply read and write the data to the optical network as if it were one large, shared disk drive. Because the intrinsic carrying capacity of a multi-wavelength optical network like CANARIEís CA*net 3 is so large, it acts as a gigantic, nation-wide optical storage device. Since all the participating computers will have ready access to all the data circulating on the network, their collective ability to solve problems quickly will be greatly enhanced.

This revolutionary use of optical networks has immediate applications in such collaborative research fields as environmental modelling, genomic and pharmaceutical simulations and astrophysics, to name a few.

"This is an exciting concept," says Andy

Continued on page 24

Chicago Civic Network - Fiber to Link More Than 1600 City Institutions - November

RFI for Condo Style Build Yields 63 Responses

Editor's Note: Joel **Mambretti** is Director of the International Center for Advanced Internet Research (iCAIR) at Northwestern University, which was established by Northwestern University, IBM, Cisco Systems and Ameritech. It has since been joined by other corporations. iCAIR works in partnership with the national and international advanced networking community. He is also Director of the Metropolitan Research and Education Network, (MREN) which is a seven-state, upper Midwest, regional network, which is run as a consortium by universities and national labs in that area. Prior to establishing iCAIR, he was Director of Advanced and High Performance Computing at the University of Chicago. He was also Director of the University's Academic Information Technologies, and Office of Strategic Technologies. We interviewed Joel on November 28 in Montreal.

COOK Report: How do you account for the emergence of the CivicNet project?

Mambretti: The primary reason is Mayor Daley's decision to developed a technology plan for the City of Chicago. However it also has been influenced by the area's advanced networking initiatives, such as iCAIR. In part, our networking initiatives in the Midwest are motivated by astrong research community that requires advanced networking. The first major iteration the implementations of the Midwest ARPA nodes, which was followed by a regional network that was created in the '80s, called CICnet. CICnetwas a mid-tier regional that linked the national NSFnet to the same universities and national labs When NSFnet was turned over to the commercial sector, the CICnet Board of Directors debated about its future, but determined a different networking model was required. We finally decided to phase out CICnet by privatizing it in the early '90s. At the same time there were some people who were saying that the commercial Internet would take care of advanced research needs.

However, some of us in the region did modeling of the requirements for the advanced research projects in the mid to late 90s and determined that we needed our own specialized, advanced networks to support those projects. So we designed MREN in 1993, we implemented it in '94. We helped create the architecture for the NAP (Network Access Point), which was established in Chicago.

COOK Report: This is called the Ameritech NAP now?

Mambretti: Yes. It was done in partnership with Ameritech, correct. It was established in '94 and it's become the basis for a number of advanced networking projects in the region and elsewhere. It turned out to be a very good model. However, that was seven years ago and now we're working on the next generation model, based upon optical networking concepts. The Chicago NAP now is the world's largest Internet exchange. So it has grown from those original concepts into a substantial facility. There are approximately 130 ISPs connected at the NAP. The history of the NAP and MREN is noted as a case study in a book that I recently coauthored with Andrew Schmidt – The Next Generation Internet by Wiley Press.

COOK Report: And the fabric it uses?

Mambretti: It's IP over ATM over SONET over fiber, which is a classic model. MREN is full mesh of PVCs and which allow very easy peering arrangements. We have also provide peering with the national research networks, such as Abilene and the vBNS. We also interchange traffic and peer with the national agency networks. Tom DeFanti, Director of the UIC Electronic Visualization Laboratory, has established at the NAP an exchange called the NSF-funded Science Technology and Research Transit Access Point, STAR TAP, which provides for connectivity among international advanced research networks and interconnects them with US national and regional research networks..

The MREN community includes NCSA; Argonne National Laboratory; Northwestern; Fermi National Accelerator Lab; University of Illinois at Chicago, the Uof I at Champagne Urbana; University of Chicago, Iowa, Iowa State, MERIT, University of Michigan, Michigan State University, Indiana University, Purdue, OARnet (Ohio State), University of Wisconsin at Madison and at Milwaukee, Notre Dame, and the University of Minnesota.CERN and CANARIE are also members, and we work closely with APAN and SURFnet.

Why Chicago

Given that the Midwest has a strong community of people involved in networking and people who design and develop such networks, and who work together well. It's a good, cooperative consortium community, which has developed some significant infrastructures. Argonne National Laboratory is a key participant. For example, Linda Winkler, who is MREN's Technical Director is a member of ANL's Math and Computer Science Division. Because of this history of cooperative partnership, it seemed only natural the iCAIR would also work in partnership with other entities, such as the State of Illinois with its Illinois Century network and the City of Chicago to help them meet their needs in advanced digital communications. Consequently, iCAIR is working with Chicago on the CivicNet Project.

COOK Report: Can you tell me how that got started and who's behind it and what the basic components are?

Mambretti: The background to Civic Net began a couple of years ago when Mayor Daley decided to develop a technology plan for the City of Chicago. It became a strategic plan for the City, to prepare it for the digital economy of the 21st century. He pulled together a team of technology advisors, called it the "Mayor's Council of Technology Advisors," and established a series of projects under that committee, including projects to address a wide range of policy issues, including economic development. Another project is working on bringing venture capital into the area. Others are focused on addressing e-government services, the digital divide, education, which is a high priority for the Mayor, and other initiatives.

The Council established a series of major initiatives, one of which I'm chairing, the Information Technology Infrastructure Committee. Under that committee, we have a variety of different activities, only one of which is CivicNet. There's many, many different initiatives, many people involved, and many committees. The people in the city have been excellent to work with on all of

these issues. It is important to understand that CiviCnet is being developed within this broader policy context. It is not an isolated project.

COOK Report: How are they being funded, primarily? Out of tax funds? Out of industrial contributions?

Mambretti: The City people are being funded by existing City budgets. There are others who are working with the City who are funded by their own organizations. Essentially, the MCTA initiative is a public/private partnership.

COOK Report: And is part of what's going on in the telecom area have some people sat down with the city folk and said, if your telecom budget traditionally for the next five years is X million dollars, if you start to invest a little of that upfront in the following ways, you can leverage what you can achieve? You can build an infrastructure that will do ten times as much for maybe even less total expenditure?

Mambretti: Well, you are bringing up many concepts in your question, so I'll take them one at a time. One is, the City does do forward planning. So they look to the future, they determine their needs, they determine what their digital communications requirements are and they do modeling that maps the demands of the future for the various agencies to budgets. So they do budgeting, they do long-term strategic planning. In looking at all that, the City is doing what any large organization should naturally do, that is to say, it is planning for a future rise in demands, and it is planning for new services, while understanding that it has a set budget.

The projected CivicNet budget is $250 million over 10 years. The question is: How can we best optimize this expenditure? Part of this process is simply doing what organizations should normally do. Certainly, then, they want to ensure that their requirements are clearly communicated to potential providers of those services, because the providers are very eager to know what customers require so that they can respond. Another part of this process is establishing an ongoing dialogue between the City and potential providers of services to match what is being asked for with what can be provided. This is a process that is healthy for both sides and one that both sides appreciate

Therefore, the idea isn't to go out say: Build something. But rather to say to the general world: Here are the requirements, and then ask for a response. That is why the city has issued an RFI not an RFP. The RFI is there to ensure that the appropriate communication process takes place.

COOK Report: An RFI, request for information?

Mambretti: Right. In part, the request for information is a request for dialogue about various possibilities.

COOK Report: Is the idea that an RFP would follow the RFI?

Mambretti: Oh, absolutely, an RFP will follow. But an RFI allows for additional dialogue and communication, discussion and research to ensure that the RFP is an optimal RFP. They are soliciting additional partners in what will ultimately become a multidimensional public-private partnership.

City Infrastructure and the Bidding Process

COOK Report: Okay, what is the situation now, then, with fiber infrastructure in the City of Chicago and suburbs that can be used by municipal government? How much exists?

Mambretti: There is already substantial fiber in the central Chicago Loop area, which is a commercial district. There is less fiber in the neighborhoods. That's a general description.

The city itself has a number of agencies that have undertaken fiber builds for their own purposes. For example, the Transportation Department implements fiber that it needs to accomplish tasks like gathering information from traffic signals. This type of fiber is already in the ground.

In addition, the City does a considerable amount of underground development. They develop for a variety of underground utility systems and they are interested in utilizing those resources for this project. In addition, in the 19th century, a complex of tunnels were built under Chicago to provide for movement of freight and mail and other types of goods, including coal. Those tunnels still exist. They form an existing infrastructure that could possibly be used. This relates to emerging new thinking related to policy development for the city. For example, everyone has recognized that continually trenching streets is a bad idea. So one of the policy aspects of this program is to minimize street trenching by having all parties involved work closely together and perhaps take advantage of some of the existing infrastructure, such as those tunnels.

COOK Report: Presumably there's lots of access to the tunnels from major buildings.

Mambretti: Yes, those tunnels go into basements of many major buildings in the commercial district.

COOK Report: So the thought is that anyone building a basic proposal would presumably use the tunnel infrastructure for the major backbone or loop or whatever the architecture would be.

Mambretti: If one is doing a fiber build under this program, there is the potential for using those tunnels to do it. That has not been done systematically in the past, and that's one reason for considering them as resources for this project.

COOK Report: And the reason that it hasn't is because the tunnels belong to the city, not to private entities?

Mambretti: That's correct. However, to some degree the tunnels have been used for communications conduits, just not to the extent that may be possible in the future.

COOK Report: Have people been preparing some estimates of what type of infrastructure you could lay using the tunnels, as opposed to traditional trenching methodologies. That would be a much more powerful outcome.

Mambretti: That is one of the planning considerations that has been undertaken under this project,, which is to say that these tunnels could be used in that fashion. Again, it has not been done systematically in the past, but under this program, there will be more of that being done. The City also owns conduit which may be used for this project.

COOK Report: What kind of schedule is being set?

Mambretti: Well, the RFI was sent out in early November 2000. The responses will be returned in January 2001. From those responses, an RFP will be developed and probably issued in the Spring of 2001. Those will be evaluated over the Summer and there probably will be some iteration of those responses and probably a contract awarded in the late Fall of 2001.

COOK Report: Who would be possible winners of a contract like that? Would it be somebody like a Bechtel? Who is your potential bidder pool? What does it look like?

Mambretti: Anyone who could provide a full range of advanced services. Anybody who could effectively meet the city require-

ments, and there are a variety of entities that could do that. What we have found is quite a bit of enthusiasm on the part of potential providers for this project. Because they see it as an opportunity to provide advanced services in ways that are now emerging that will become commonplace in the future.

COOK Report: Are you talking to some of the inter exchange carriers? Are you talking to the ILECs?

Mambretti: This project actually has attracted attention from numerous entities. A meeting was organized as a follow-up to the release of the RFI that attracted about 200 attendees from many companies. I don't have a complete listing of respondents, the city purchasing office is actually keeping a record, so for that information, you might want to go to them and see who is responding.

COOK Report: Again, I guess no one is attempting to prejudge what the overall structure would look like, but what are some of the potential components of this? Are you likely to have some people that are just specialists in laying fiber? How do you decide what kind of architecture to put on your fiber? Is it likely that it'll be gigabit Ethernet because of cost considerations? What about people who say, oh, my god, we need SONET, we need ATM and all that stuff?

Technology Assumptions

Mambretti: Well, there are some parameters that are very clear. For example, the RFI itself states that Gigabit Ethernet will be an important service that will be provided through this system, as well as 10G E. That's in the spec. They will have to provide those services, as well as other high performance services.

Also, it says that the system has to provide gateways to traditional systems, such as ATM and SONET. But given that it says there should be gateways to those systems, it also implies that perhaps the core backbone will be based on current emerging optical technologies, as opposed to those traditional technologies. Taking account of the fact that digital communications are moving towards a model which is powerful and simple in the core and has complexities at the edges, as opposed to today's ITU model, which is complex in the core and simple at the edges. This architecture will take that trend into consideration. The architecture will also reflect many IETF concepts, not just ITU concepts.

COOK Report: And what are some of the IETF concepts?

Mambretti: Well, I just named one. IETF architecture is based on a premise of having a much simpler core and complexity at the edges of a network, as opposed to the other way around. Another is having multiple paths supporting connectionless digital architectures as opposed to SONET rings supporting a circuit-switched, connection oriented architecture. And in part that helps reliability through redundancy by avoiding a single point of failure.

COOK Report: So if one were to guess, I suppose what this basic core would look like, it would be a ten gigabit infrastructure, into which city agencies or commercial entities or anybody could plug whatever they wanted to do. Other than city agencies, public schools and so on, what has been the discussion of where the borders lie between commercial telecommunication entities and what the city is doing?

Mambretti: The city does not intend to get into the provider business. The city would prefer not to, as a matter of policy. Therefore, what the city is requesting is a consideration of the city as kind of an anchor tenant within a new type of system, where others could join into that consortium. So there is an anticipation that others will join and that what will be developed will be a core backbone doing not just gigabits, but more than tens of gigabits, more than hundreds of gigabits in the core backbone. You will have a very strong core backbone infrastructure that links to a very wide variety of tributary services, including links to traditional services, such as ATM, and SONET. But it will also be broad enough to include such things as broadband wireless and other types of services.

So the idea then is to make the whole system more flexible, more adaptable, more customizable so that various people can plug into it in a very easy way.

COOK Report: Optical technology right now would be understood, would it not, more as Layer 2, but what type of Layer 3 IP routing considerations are a part of this?

Mambretti: The architecture provides for standard Layer 3 routing, supported by new techniques and technologies that support IP over optic infrastructure, including signaling mechanisms such as GMPLS. Increasing, IP Layer 3 techniques will be used for service provisioning, traffic management, and even optical layer function control.

COOK Report: And supported by the optics underneath.

Mambretti: Correct. Provisioning will be accomplished through light waves, for example, setting up, deleting, and swapping individual light paths. It is possible to provide for multi-service provisioning by dedicating individual light paths to specific services.

COOK Report: Right, lambdas. Do you look at OBGP?

Mambretti: That's also mentioned in the RFI. CANARIE and Bill St. Arnaud, have been great inspirations for our efforts on optical networking, especially for metro optical nets. We are examining various IP-over-optics architectural models, including overlay, signaled overlay, peering, and integration.

COOK Report: As a means of very high bandwidth routed connection from point A to point B?

Mambretti: Yes and also to assist with dynamic provisioning, routing traffic management, reliability and survivability guarantees, performance and overall network management.

Question on February 10, 2001

COOK Report: Please summarize what response were received by the January 19, 2001 RFI deadline.

Mambretti: We are pleased by the enthusiastic response to the RFI. Approximately 63 individual organizations and consortia responded. The directions outlined in the RFI related to future digital communication services and technology architecture have been widely endorsed. This initiative also conforms with traditional City policy, which has often been oriented to enabling the economy through infrastructure development. In the past, this development has been oriented to transportation and utility infrastructure. It is natural to extend this activity to enabling the digital economy through the development of advanced communications infrastructure. Over the next few months we will be gathering a broad range of additional information as preparation to developing the RFP, and our expectation is still that the final document will be issued in late April or early May. All participants in this particularly innovative project have appreciated the opportunity to assist in making a great City even better, through advanced technology. Eventually, all cities will become digital cities, illuminated by light waves.

Editor's Note: Thanks to subscriber Frank Coluccio we received a pointer to the Civic Net web site that has been built by the city of Chicago.
http://www.chicagocivicnet.com/civicnet/SilverStream/Pages/civiclist.html The site and the RFI document itself – some 130 pages in Adobe PDF format - begins to bring home the seriousness of the project. For a project of this scope, it seems likely that the city is making unprecedented use of the Internet to bring a public focus to its civic net acquisition procedures. In addition to the RFI the web site contains a forum for respondents to question the city and it contains access to city mapping tools and map sets that in non electronic form would be hugely expensive to enlist in a venture such as this. Finally it is perhaps the most significant example of CANARIE's telecom outlook influencing events in the United States. We republish here a shortened form of the RFI text that is also lacking its extensive appendices. The complete document is worth examining in order to get a sense of how ambitious this project is.

Request For Information Chicago CivicNet

Specification No. B09189503
Required for use by:
City of Chicago Department of General Services ,Bureau of Telecommunications and Information Technology

This RFI distributed by:
City of Chicago (Department of Purchases, Contracts, and Supplies)

All Responses must be addressed and returned to:
City of Chicago, Department of Purchases, Contracts, and Supplies Bid and Bond Room -Room 301 City Hall 121 North LaSalle Street Chicago, Illinois 60602

1.0 INTRODUCTION

1.2.0 There are five (5) general objectives that the City would like to see addressed in implementing this infrastructure. First, this initiative addresses needs for a broad range of required services for multiple entities and constituencies, as described below. Second, it is oriented toward emerging communication capabilities rather than current traditional services and technologies. In part, this initiative addresses long-term strategic goals. Third, it presumes the development of new, rather than traditional, business models for service development and support. Fourth, the City intends to expand this initiative to include additional partners, in the public and private sectors. Finally, related to these new business models, a fifth objective is to formulate, in partnership with services providers, new types of cost and pricing structures for the mutual benefit of all parties.

1.2.1 CivicNet must consider the needs of multiple constituencies in multiple locations.
1.2.1.0 The City and the Parks are currently served by a digital Centrex (CityNet) of thirty-five thousand (35,000) lines in eight hundred fifty (850) locations, some of which are switched data lines, in addition to an assortment of T-1, multi-drop, frame relay, ISDN, and other data carriers. Twelve (12) downtown locations are linked via leased SONET services or dark City-owned fiber.
1.2.1.1 The Chicago Public Schools are currently served by an analog Centrex system with eighteen thousand (18,000) lines in six hundred (600) locations, in addition to an assortment of data lines that are primarily T-1 speed lines.
1.2.1.2 The other CivicNet agencies are served by similar leased services. Detailed information on existing networks is provided in Section 5 and in (Appendix A).
1.2.1.3 These networks are experiencing growth in usage and new requirements for expanded geographic coverage, as well as significant increases in bandwidth. An equally significant increase in costs will be inevitable unless alternatives to our current networks can be developed.

1.2.2 The Chicago CivicNet is intended to address the combined needs of these existing networks, to be created by upgrading the technology used in the existing networks and expanding their geographic reach, as well as allowing for improved controls, monitoring, and billing services, all of which can take advantage of economies of scale.
1.2.2.0 The City desires proposals for all hardware, software, and services required to handle the current network demands, utilizing initially the in-place customer premises equipment, as well as to upgrade and operate this network on an ongoing basis. The City does not anticipate that one vendor will be capable of meeting our varied requirements and thus the City is hoping for a solution where multiple players, large and small, will work together to provide a unified solution.

2.0 CHICAGO CIVIC NETWORK CONCEPT

This is the Mayor's objective: "I envision the entire City -residents, businesses, and institutions -using the network to access on-line education programs, video-on-demand services, telecommuting, and on-line community organizing."

2.0.1 The following summarizes the concept being pursued:
2.0.1.0 The City owns, or has access to, resources of significant value that can be used as the foundation of a fiber infrastructure. These resources are summarized in Section 7 and further detail is provided in (Appendices B and C). Examples of resources that could be made available to the City's "partners" in development of CivicNet include the underground freight tunnels downtown, CTA right of way, City-owned fiber already in place, and available duct and conduit that runs throughout much of the City. The City can install conduit and fiber every time a street is opened to lay new water or sewer pipe, etc. As a result, the costs usually associated with building a network in an urban area will be greatly reduced.
2.0.1.1 The City agencies constitute a major user of telecommunications services. Their current combined expenditures, exceeding $25 million annually, are purchasing leased services. This $25+ million could be likened to "rent." The City would prefer to be in a position of building equity, with the eventual result being lower or stabilized costs, and at some point a revenue source for the City. Most importantly, CivicNet is seen as a significant boost to Chicago's continued economic development.
2.0.1.2 The City, itself, does not want to be in the business of marketing, wholly financing, and operating a communications company. The City prefers to leave those responsibilities to the experts. However the City does intend to work with communication wholesalers and service providers to insure that its specific requirements are well understood and that those requirements are met. The City will not simply accept generic solutions that do not meet its explicit requirements. The City wishes to benefit from being in an equity position, in a relationship with multiple entities, possibly organized into a public/ private consortium, on contract with us to build, market, operate, and manage the network infrastructure.

2.1 Service Guidelines The following guidelines shall govern the service structure of the proposed network infrastructure:
2.1.0 Equal access -The City expects equal access throughout the network. All informa-

tion content and service providers requesting access to the network infrastructure shall be treated equally. Facilities, wholesale, and retail services should be made available in a fair, reasonable, and nondiscriminatory fashion.
2.1.1 Open platform -Technical compatibility and interoperability among network providers must be assured. The network must be standards-based, capable of being connected to other private and public networks.

2.1.2 Abundant bandwidth -There must be available capacity beyond current demand, with the vendor responsible for building in the flexibility required for incremental growth as market forces dictate. Irrevocable rights to use (IRUs) should be made available on a reasonable, fair, and nondiscriminatory basis.
2.1.3 Network availability -CivicNet must have the capability to connect to any location throughout the geographical area of the City. Because the City agencies are located in every neighborhood, this becomes a non-issue.

2.2.4 The twenty-five million dollars in annual expenditures that the City agencies spend for current services might place the City in the category of anchor tenant for the development of CivicNet. Revenues from the other users of the infrastructure from the institutional and private sectors will likewise be directed to the winning vendor(s) who partner with the City in this endeavor. To work, this must be a win-win for all parties involved.

2.2.5 The City expects to realize the following goals by implementing the telecommunications infrastructure:
2.2.5.0 Minimize the adverse financial impact of technological service upgrades on the City agencies.
2.2.5.1 Ensure that the City receives fair compensation for the use of public property as appropriate under state and federal telecommunications law.

2.2.5.2 Provide for the optimal use of public right-of-way, avoiding redundant infrastructure, and minimizing disruption to property and the public streets by providing a common, universally accessible telecommunications infrastructure, with the understanding that nothing herein shall prevent the use of the City's public ways by third parties consistent with and subject to applicable City regulation.

2.2.5.3 Ensure that residents, businesses, and institutions will have universal access to reliable, high-speed, high-performance telecommunications services at the lowest feasible, competitive cost via an optical network infrastructure. The City wishes to encourage use of the public rights of ways toward this end.
2.2.5.4 Encourage economic development by stimulating the growth of the regional economy and enabling businesses to effectively participate in the global economy.

2.2.5.5 Ensure that the maximum number and variety of telecommunications services are made available through effective competition.

3.0 BUSINESS CASE

3.1.3 The City agencies currently have sixteen hundred (1,600) locations that need to be connected in the reasonably near future to a higher-speed network than now exists. Unfortunately, there is no train track running to each of these facilities.
3.1.4 So the challenge is: Can a business case be made that justifies the expense of building an infrastructure to service every part of the City of Chicago
3.1.5 For the bottom line to make sense, two things have to work together: The cost of building the infrastructure has to be lowered, and, the opportunity to bring in revenue has to be enhanced.

3.2 General Concept
3.2.0 In layman's terms, it is this:
3.2.1 The City agencies are spread throughout our geographical boundaries. The City spends in excess of $25 million every year on leased services. If the City were to build out a network infrastructure all over the City for our own purposes, it would make sense to let the private sector use that same infrastructure. Merely by connecting all (1,600) of the City's properties to CivicNet, the network infrastructure would extend to encompass three million City residents and businesses. The City also owns substantial resources that will reduce the cost to build this infrastructure. So, who wants to partner with the City in building a network infrastructure for all of its constituents to use? The City's leased services could fill the role of anchor tenant in the project, and the extension of the CivicNet infrastructure to the private sector could provide a significantly enhanced revenue opportunity for CivicNet partner/ providers.

3.2.2 There are a number of roles that different entities might choose to play. Those entities that contribute resources to help make the network a reality would be rewarded accordingly. This could include:
3.2.2.0 Providing access to the City with a low-cost means for network construction, possibly in exchange for fiber or other considerations.
3.2.2.1 Providing capital to help finance the project in exchange for long-term use of fiber to selected locations. Examples: Institutions with branches scattered around the City that want to connect them with fiber they control (banks, chain stores, hospital and medical groups, universities).
3.2.2.3 Providing capital to help finance the project in exchange for long-term use of fiber throughout the area. Example: Companies competing for a share of the telecommunications industry. Examples of these companies include service providers, carriers, carriers' carriers, CLECs, etc.
3.2.2.4 Providing venture capital in exchange for future revenues.

3.3 The Opportunity The City owns resources that can provide a new and lower cost approach to building high-performance infrastructure that has wide application throughout a major metropolitan area.

3.3.0 Lowering the Cost of the Build Section 7 contains a description of many of the existing City resources that can be used to build out the CivicNet fiber-based infrastructure. These include:

3.3.0.0 City-owned fiber that connects about one hundred (100) City agency locations.
3.3.0.1 City-owned fiber that totals several thousand miles, most of it single mode, most of it in prime locations. Recently, fiber has been installed to reach new Police locations and Department of Transportation locations throughout the city.
3.3.0.2 The Chicago Transit Authority has extensive rights-of-way. (Appendix B1)
3.3.0.3 The extensive network of downtown underground freight tunnels. (Appendix B2)
3.3.0.4 The extensive network of four-foot diameter (and larger) storm sewer ducts that are under every neighborhood in the City and are interconnected. (Appendix B3)

3.3.0.5 City franchises with utilities and other companies that allow for City use of existing resources, to the extent available and where such use by third parties for governmental and other purposes is permitted.
3.3.0.6 City control of the water, sewer, street light, and transportation functions, allowing for conduit to be installed whenever new water pipe, sewer pipe, street improvements, etc. are undertaken. This construction totals hundreds of miles per year ^ every year. NOTE: Over one hundred (100) of the largest City locations are already connected by fiber in varying quantities that is not being fully used and is available.

3.3.1.5 The following example outlines one possible approach to making CivicNet a reality:
3.3.1.5.0 Partnership (Consortium) XYZ

Corporation proposes a Consortium approach to the basic fiber optic backbone that will be composed of both public and private corporations. The Consortium will be organized as a (not for profit/ profit oriented) organization to meet the needs of the Consortium members. Organizations that might choose to join the Consortium could include:

Private Companies: Local Exchange Carriers Competitive Local Exchange Carriers Carriers' carriers Cable TV companies Equipment suppliers Contractors Utilities Private companies Healthcare organizations Institutions: Universities Hospitals Not-for-profits Museums Governmental entities Corporations

3.3.1.5.1 Managing Agency The XYZ Corporation will be the managing partner of the Consortium with the following responsibilities:

Fiber Structure Market and promote the consortium. Target markets being proposed for economic development. Plan the implementation process. Build the fiber optic backbone structure, utilizing City resources wherever feasible. Building the structure will include construction, installation, cabling, utilizing varied means of access, etc.

Provide regular maintenance that promotes the reliable operability of the system. Maintenance will include network, node, and network operation center functionality. Provide monitoring of the fiber optic system. Provide an architectural design and outage restoration response that promotes 99.999% reliability within the system.

Municipal Applications (more "sample" language such as might be provided in the eventual solution) XYZ Corporation has partnered with ABC Corporation to meet the specific needs of the municipal organizations as outlined within this RFI. Attached are proposals that address the building and operating of a fiber optic infrastructure and providing the value added services for the City agencies.

3.3.1.5.2 Public/ Private Consortium The Consortium recognizes that various partners have contributed to CivicNet either monetary or material resources. The following generic financial model illustrates a concept of how the business case can be actualized: (Responders to provide) Assumptions System Buildout Maintenance Network Operations Center Outage Restoral Capacity Marketing and Growth Market Potential

3.3.2 Varying Levels of Participation are Encouraged 3.3.2.0 The following are examples of the range of responses the City anticipates receiving to this RFI: 12

13 3.3.2.0.0 Providing access to the City/ CivicNet with a low-cost means to build out the network construction, possibly in exchange for fiber or other considerations. This could apply to utility-type companies, among others.
3.3.3.0.1 Parties interested in purchasing "condominium" fiber to run to connect selected locations, for the parties' exclusive use. This could apply to financial institutions, chain stores, universities, healthcare organizations, technology companies, etc.
3.3.3.0.2 Providing capital to help finance the project in exchange for long-term use of fiber to selected locations. This could apply to the same entities as above.

3.3.3.0.3 Providing capital to help finance the project in exchange for long-term rights to use fiber throughout the area. This could apply to service providers, LECs, CLECs, etc. 3.3.3.0.4 Providing venture capital in exchange for future revenues.
3.3.3.0.5 Parties interested in leasing fiber and/ or services carried over fiber.
3.3.3.0.6 Companies interested in supplying certain elements of technology to the project. This could apply to equipment companies, possibly interested in partnering with other entities.

3.3.3.0.7 Service providers interested in supplying services to the City and/ or to the private sector.
3.3.3.0.8 Systems integrators interested in managing all or major pieces of the project.
3.3.3.0.9 Construction companies interested in the build-out of the fiber.

4.0 REGISTERING FOR THE RFI

4.0. All prospective CivicNet RFI respondents should complete the mandatory online registration form found at the World Wide Web (WWW) address indicated below:

4.0.1.0 www. chicagocivicnet. net.
4.0.2 By completing the online registration form, each RFI registrant will:
4.0.2.0 Establish itself as a Company of Record for City of Chicago documentation purposes;
4.0.2.1 Gain access to the electronic version of the CivicNet RFI;
4.0.2.2 Gain access to the City of Chicago's Geographic Information System (GIS) maps, which identify the locations of the 1600 City agency's CivicNet sites; (Note: This electronic file is NOT available in the paper-based RFI document.)
4.0.2.3 Gain access to the Web-based CivicNet RFI Question and Response board. (See Section 10.2)

4.0.3 Any respondent that experiences difficulty accessing any of these CivicNet RFI Web-based files or has questions regarding this process, should contact the City by e-mail at civicnetrfi@ cityofchicago. org.
4.0.3.0 Questions related to process and procedure can also be called in to 312-74CIVIC (312-742-4842). Questions concerning the specific contents of this RFI will not be answered over the phone. Content questions must be addressed to the web site.

Continued from page 18

Woodsworth, Director General, Institute for Information Technology, National Research Council of Canada. "We operate Canada's largest distributed network of servers as part of The Canadian Bioinformatics Resource, a national service dedicated to providing researchers with convenient, effective access to high-quality genomics and bioinformatics tools and databases, interconnected to our new 6912 processor Paracel GeneMatcher in Halifax. This technology will allow more efficient distribution of data between these servers."

The implications of this concept to telecommunications carriers could be equally significant. Purchasing a super computer in the future might simply involve buying optical wavelengths from a favourite carrier just as silicon chip based machines are purchased from computer manufacturers today. This revolutionary use of optical network technology could also become an important use for unused bandwidth. No longer would networks be simply a means of computers exchanging data directly with each another; the network will be the computer.

CANARIE will be deploying a proof-of-concept wavelength optical disk drive on its CA*net 3 network over the next few months.

Congress Gives ICANN Second Look

Auerbach and Froomkin Testify Before Senate that Compared to House is in Early Learning Stage

Editor's Note: In early February the House Commerce Committee held hearings on the selection process that ICANN used in deciding on new top level domains. They were followed less than a week later by hearing before the Senate on February 14.

Having listened to the hearings with real audio, we find that compared to the hearings of the preceding week, the February 14th hearings were an anti climax. The supreme moment came when Senator Burns asked Michael Froomkin what he would propose in light of the criticisms of ICANN that were delivered to the Committee.

Froomkin: "the most important issue is not setting a precedent by which a Department, like the Department of Commerce can end run the Administrative Procedures Act. And that is an issue that frankly is bigger than the Internet. The global concern here is not just in this process. For this is a way in which agencies can by pass ordinary procedures to create a privately organized regulator in all but name. A regulator that uses control over a federally dominated resource to make people sign contracts with it, pay it money and do what it says. And then not be subject to due process. Not be subject to court challenge. Not be subject to ordinary oversight. That is really cutting the congress and cutting the American people out of the regulatory process. So while in the case [of the seven selected new TLDs] you might have had an outcome which was better than no decision at all – I have nothing against any of the winners here. I have no reason to believe that any are bad or imperfect, and for all I know we would be all better off if they were all put in the root along with lots of others too. It seems to me that there is a good government issue that is pretty serious here. Someone needs to hold Commerce's feet to the fire on this one."

Thirty minutes later at the end of the hearings, Burns: "I think we are going to have another hearing because there are questions out there yet and once we dive through your information and the testimony that was held here today. I am concerned about redress of due process. I am going to go through this MOU so that I am a little more enlightened on the situation between commerce … was that transfer made…. Was it made legal? Maybe we might do it thataway? Not necessarily getting into the operation of ICANN…. But let's make sure that the parties understand….. Let's lead out a little bit better where there is due process and some of those things that I think are important in the business world. Before the technology can reach its full potential, the users have to be confident that they are being protected." Boxer, whose performance as shill for the Intellectual Property interests was especially abysmal said that she also agreed.

The Senate Doesn't Get It

The Senators were not well informed and while Auerbach's and Froomkin's testimony was accurate and should have caused rapt attention, Mike Roberts and Roger Cochetti as two of the people most responsible for the mess sat there and made apologies for an ICANN that really was not yet mature and had not yet had a chance to do what it was put into place for. They and fellow insiders from the .cc and .edu registries, being insiders to the process, gushed about how free and open it was. Cochetti was happy about ICANN as was Cartnell from the .cc registry. Both should be happy because both are in the position of making vast sums of money from the DNS system as currently operated. Unfortunately it is highly unlikely that Burns or Boxer (unless they have already been 'bought') have a clue about the real reasons for the state of mind of those who defended ICANN. Without understanding the details of the ICANN jungle, (Burns said twice his eyes were glazing over) it is unlikely that either will ever be motivated to take any constructive action.

Michael Froomkin offers the only reason for any hope. Here we have an attorney who after a stint in the London office of Wilmer Cutler joined the faculty of the University of Miami Law School. A specialist in administrative law Froomkin became interested in the potential of the Internet to empower the individual citizen. With no real experience in the domain name wars, he decided to attend the First of the IFWP forums in Reston Virginia on June 30 – July 1998. He became hooked and for the next 30 months devoted vast amounts of time to researching and writing about ICANN from every parameter. In December 2000, the Duke University Law Journal published his 168 page legal analysis: Wrong Turn in Cyberspace – the definitive legal analysis of this quasi-governmental organization's mission to take advantage of American distrust of government and enable private corporate control and regulation of the Internet. ICANN is performing its mission with none of the administrative and due process procedural safe guards to which citizens would be entitled had the job been left in the first place for the US government to accomplish. Corporate supporters have plans for many ICANN-like organizations to follow. Froomkin's work shows this clearly.

To his great credit whenever the question of ICANN's legal legitimacy comes up Froomkin is now considered to be the primary expert. The problem is that for his testimony to swing the tide, he needs serious allies and many more adherents. In the meantime Roger Cochettti the IBM installed satrap of Network Solutions will continue to tell Congress that NSI-Verisign really supports ICANN and has proven its good will by contributing more than anyone else to its budget. Congress will never understand the quid po quo that NSI-Verisign does so because under ICANN it has been able to prolong and extend the monopoly it enjoys – the very one that ICANN pledged to undertake a holy mission to eliminate.

Auerbach Versus Roberts

The hearings of Februray 14th set up a stark confrontation between ICANN President and insurgent Board member Karl Auerbach and outgoing president Mike Roberts.

Herewith some comments on how this shoot out came to be.

On Mon, 12 Feb 2001 On C Y B E R T E L E C O M - L@LISTSERV.AOL.COM in response to the announcement of the Senate Panel Steve Wolff wrote: Go Karl!

Dave Hughes responded: Ah yes. It seems, now that there is a new Administration in town, the whole question of ICANN - its genesis, support by the *old* administration (Ira Magaziner, then a murky corner inside the Department of Commerce, and the NTIA) - is going to be vetted by the new Administration, starting with Republican Congressmen. Maybe we will all find out, at last, what and who has *really* been behind ICANN, which itself is a front for other interests. Not to speak of 'entangling alliances' with other nations and representative organizations like the ITU.

And people actually think ICANN is no more than an Internet 'name' assignment body. Rather than a badly concealed mechanism which eventually could/would 'control'

(aka for 'governance') the Internet, on behalf of both private, and public, interests.

We added: Oh a *lot* of people are catching on Dave to what is going on. I commend Steve Wolf for cheering on Karl Auerbach who while a Cisco employee ran as insurgent last summer. Using the long odds of ICANNs self nomination system KARL beat out 4 other ICANN nominees including Lawrence lessig and two other self nominees to win the at large members seat for North America on the ICANN board. He ran in no small part because Mike Roberts ICANN president had usurped the organization on behalf of the interests of IBM and other intellectual property holders. Roberts was leading an organization that pretended to be founded on consensus. But Under Roberts its behavior was secretive, authoritarian and devious. Roberts, certain that he knew all the answers, plunged arrogantly ahead while dismissing his critics as arrogant school children. Auerbach ran and was elected on a platform that called for Robert's removal.

Roberts, long time friend of IBM funded Educause, had been working behind the scenes in the summer of 1998 to ensure an Internet society take over of the the nascent "new co" on behalf of the very people who had piloted the defunct IAHC gtld effort. Much of the summer had gone by with don Heath's men wondering when and how to overwhelm the International Forum on the White paper process. While at the same time Joe Sims, under premises that turned out to be without foundation, convinced Larry Lessig to enter the negotiations. Lessig, friends say, never forgot the deception. In mid August Roberts horse traded in the background to break up plans for the wrap up meeting of the IFWP peeling allies off one by one to destroy any chance that an organization other than the Jones day reavis and pogue designed icann rump would emerge with any legitimacy.

On August 28 in the imperious style he would soon adopt as Icann leader Roberts announced his withdrawl from the IFWP process.

Dave Farber took Roberts note with permission and lobbed it bomb like into his 25,000 person Internet gatekeepers mail list AKA "IP"...whereby with this action he effectively told the key opinion shapers that the only non ISOC alternative to ICANN was dead.

From: Dave Farber <farber@cis.upenn.edu>
Subject: IP: "Ratification" - the IFWP Emperor has no Clothes

The following is sent WITH permission. For those of you who don't know Mike, no one would ever characterize him as a flamer or a radical. It is worth reading and thinking about.

Dave

From: Mike Roberts
<roberts@ivory.educom.edu>
Subject: "Ratification" - the IFWP Emperor has no Clothes (fwd)
To: farber@cis.upenn.edu
Date: Fri, 28 Aug 1998 17:40:35 -0400 (EDT)

Forwarded message:
From: Mike Roberts
<roberts@ivory.educom.edu>
Subject: "Ratification" - the IFWP Emperor has no Clothes
To: ifwp-discuss@itu.int
Date: Fri, 28 Aug 1998 15:06:21 -0400 (EDT)

I'm appalled at the convoluted and narcissistic thinking behind the so-called "ratification" meeting. Let's review the bidding, shall we?

- we'll hold the meeting immediately after the yet-to-be-confirmed final-negotiating meeting of the yet-to-be-determined number of important White Paper stakeholders that we have 100% conviction will lead to an accepted legal structure for the new IANA corp.

- even though the cumulative probability of the previous item is in the range of 30%, we'll find some gullible soul with a high limit credit card to start signing financial commitments for the meeting even though we haven't come close to paying our bills for previous events

- the close conjunction of the ratification meeting with the final negotiating meeting is so tight that no one outside walking or driving distance of the site will attend on less than a full fare airline ticket and a full fare at a fully booked hotel. This is designed to ensure an open and inclusive participation in the meeting.

- also designed to ensure the broadest participation of all concerned Internet citizens, we will arrange for a electronic vote on the final document, using a system whose features are unknown to essentially all of the proposed participants. In the interests of time, we will forego any attempt at verification of the votes with actual living humans and we will not be overly concerned if the agreed upon final document is neither sanctioned by governing boards of the parties to it in time for the vote, nor that the vote commences before it is even posted.

- in view of the fact that the hosts of the negotiating session and several of the known major players have firmly rejected the notion that the agreed upon document is actually open to revision by the Internet community at the ratification meeting, we will arrange for moderators to paper over any dissents that emerge. The Chair of the meeting will call for a voice vote and declare that he/she hears more ayes than nays, thus sanctioning the decision reached before the meeting.

Thanks but no thanks. My organization will have nothing to do with such an unrepresentaive, undemocratic, uninclusive gathering.

- Mike Roberts
EDUCAUSE

Cook Speaking: pride goes before the fall..... If we look at the style that Roberts adopted as ICANN president, here is what Roberts is saying. "This is my leadership style. The only thing that bothers me is that i am not in a position to do the leading. "

In late October shortly after Roberts victory was announced, a source close to Roberts wrote: in n the spring and summer of 1998 "Mike was representing himself as a representative of EDUCOM, however, Mike never once sent out any update/report or queried the EDUCOM membership on their point of view." The source had misgivings about what lay ahead: based on knowing the man for well over a decade " 1) He feints listening and will do whatever he pleases. 2) He views his "solution" as the right answer, and does not accept criticism well. 3) Mike's "history" at EDUCOM and Internet2 has been one of little oversight. He was able to develop and carry out an agenda with little if any checks/balances from the general membership."

Cook: Those who put ICANN together did so in secret. Two and one half years later the closed undemocratic clique which thought it could fool the public by turning to names like Esther Dyson has alienated most of those on whom it should depend were it actually to attempt the technical coordination that its charter said it was there to do. Roberts is now on his way to sing his swansong in front of the senate.

However Roberts' arrogance gave rise to the very process that put Karl Auerbach in the same Senate hearing room as well on St

Continued on page 37

Hearings on ICANN Governance

Prepared Statement of Karl Auerbach before the Senate Commerce, Science and Transportation Committee

February 14, 2001

Good morning. I have been involved in the Internet since 1974 and have actively participated in the transition of its administration to the private sector for the past five years.

I am a computer engineer - I do research pertaining to ways of making the Internet more reliable and efficient in the Advanced Internet Architectures Group of Cisco Systems in San Jose, California. I am also working on a joint Cisco - University of California research project on advanced control and provisioning mechanisms for the net.

With respect to my service on the board of ICANN and for all the opinions I am expressing here today I neither represent nor speak for Cisco. My opinions are my own. I am also an attorney. I graduated cum laude in 1978 from Loyola of Los Angeles specializing in commercial, international, and administrative law. Although I maintain my status as the member of the California Bar and the Intellectual Property Section of the California Bar, I am not engaged in active practice.

I have been a founder, principal, or first employee in several Internet related start-up companies. These include Epilogue Technology Corporation (now part of Wind River Systems); Empirical Tools and Technologies Corporation; Precept Software (now part of Cisco Systems); and InterWorking Labs, Inc. These have provided me with a broad base of experience in commerce and technology. I have direct experience with the needs and obligations of Internet related businesses. I am sympathetic to the needs of Intellectual Property owners. I own copyrights, I have owned Federally registered trademarks, and I have filed for patents.

I have been active in the core design and standardization body of the Internet, the Internet Engineering Task Force (IETF), since the mid 1980's. And I have been a member of the Internet Society (ISOC) since its formation.

I have been deeply involved during the last several years with the evolution and activities of what has become ICANN. I am a founding member of the Boston Working Group, one of the groups that submitted organizational proposals to NTIA in 1998 in response to the so-called "White Paper."

I am a member of the ICANN Board of Directors. I was elected to represent North American Internet users. Four others were elected at the same time to represent other regions of the world.

I was elected to my seat. I was not appointed. I was elected to represent the Internet users of North America in an election in which I ran against six other highly qualified candidates: the Chancellor of the University System of Maryland, the Chief Scientist of BBN Technologies, the President of the Information Technology Association of America, the former President of the Association for Computing Machinery (ACM), a Professor of Business at the University of Texas, and the former holder of the Berkman Chair at the Harvard School of Law.

I am the only person on ICANN's Board of Directors who was elected by the Internet users of North America.

I have only been on ICANN's Board of Directors for a few months - my term started shortly after the well-publicized and controversial selection of a mere seven new Top Level Domains (TLDs).

However, despite the short time I have been a Director, I have already learned much to confirm my fears that ICANN is suffering from a lack of public process, lack of accountability, mission creep, poor communication, excessive delegation of policymaking to staff, and poor business practices. As a Director it is my job to work to correct these weaknesses. But I despair at the immensity of the task.

My primary focus within ICANN is on limiting ICANN's scope of authority, creating well defined procedures for fair decision making, and establishing solid business practices. These are conservative and reasonable goals. As both an engineer and an attorney with long experience in other Internet governance organizations I have a solid grasp of the issues ICANN has sought to address.

ICANN should be based on viable real-world ideas and processes, not on some abstract notion that suggests that ICANN can somehow fly above technical, economic, and political realities. Those who use the Internet ought to have a voice in the running of the Internet. I do not subscribe to the notion that people can be properly represented by pre-defined, one-size-fits-all "constituency" structures such as are found in much of ICANN's present structure.

I support the continued existence of ICANN. ICANN is a valuable institution; its roles as a technical coordinator are quite properly needed for the smooth functioning of the Internet. However, ICANN is ill designed, has been ill operated, has brought upon itself significant ill will within the Internet community, and has greatly exceeded its proper scope. I believe that significant restructuring of ICANN is needed so that the corporation can fulfill its purposes [1] and fulfill its obligations towards its stated beneficiaries. [2]

I would like to discuss the following matters: 1. What kind of entity is ICANN? 2. How can ICANN obtain acceptance and legitimacy? 3. ICANN's obligation to have meaningful public participation in decision-making. 4. Specific things Congress, the US Department of Commerce, and ICANN ought to do.

1. What kind of entity is ICANN?

ICANN is Internet governance. ICANN is far more than mere technical coordination. ICANN is about policies for the allocation of Internet resources. ICANN is responsible to no one other than the Attorney General of the State of California.

ICANN's policies have an economic impact that is potentially measured in billions of dollars. The impact of decisions in the Domain Name System (DNS) space have been well noted elsewhere and were illustrated last week in hearings before the House Subcommittee on Telecommunications.

There are those who say that ICANN is merely a technical body. I am a technologist. Yet I have a difficult time understanding how any of ICANN's decisions concerned with the Domain Name System have any technical content at all.

One must wonder where the technical component might be in ICANN's Uniform Dispute Resolution Policy - a policy that expands the protection of trademarks to an extent not granted by any national legisla-

ture. And one must also wonder where the technical component might be in ICANN's preservation, indeed in ICANN's extension, of the hegemony of Network Solutions over the naming systems of the Internet.

In other words, ICANN is very much a regulatory body. And it is a regulatory body that has been flung into existence with the support and aid of the US Department of Commerce. As if to underscore that ICANN is a fruit that has fallen not far from the administrative agency that engendered it, ICANN was created with the express purpose of "lessening the burdens of government". [1] But unlike regulatory bodies that are part of the government, ICANN is a private corporation [3] and is not obligated to undertake any of those troublesome constitutional Due Process burdens imposed on governmental administrative bodies. ICANN is not subject to the burden of judicial review or the Federal Administrative Procedures Act. ICANN is not required to make it possible for those affected by its decisions to participate in the making of those decisions. And there is no mechanism to compel a truly independent review of ICANN's actions.

ICANN's internal mechanisms for review are moribund or exist only as paper placeholders. I have had one request for independent review pending for nearly a year because ICANN has been too busy galloping off doing mission-expanding policy development, leaving it no time to pay attention to the implementation of fair procedures for review.

And until a viable at-large mechanism is created and full rights of membership accorded, ICANN has no external entity to which it is accountable other than the Attorney General of the State of California.

ICANN has gone so far as to assert in amicus briefs that ICANN believes that it, not the courts, should be the forum for resolution of disputes.

ICANN is the result of a strange brew of governmental powers and private lack of accountability.

ICANN, despite its claims to the contrary, is extremely secretive. We know more about how the College of Cardinals in Rome elects a Pope than we do about how ICANN makes its decisions. As a member of the ICANN board I have been surprised at how often I learn of ICANN actions from outside third parties. And I have perceived a very strong resistance on the part of ICANN's staff to opening its activities, even to members of ICANN's Board of Directors.

ICANN has several internal committees and organizations that have no distinct legal existence apart from ICANN. As a Director I am responsible for the assets, liabilities, and actions of these bodies. Yet some of these bodies act as completely autonomous, independent, and often very secretive entities. At least one of these entities maintains distinct financial records that seem not to be incorporated into ICANN's overall financial statements. Another refuses to allow Directors to inspect its activities or meetings.

I have a hard time reconciling ICANN's opaque processes and structures with the obligation in its bylaws that "[t]he Corporation and its subordinate entities shall operate to the maximum extent feasible in an open and transparent manner and consistent with procedures designed to ensure fairness." (ICANN By-Laws Article III, Section 1.)

ICANN has an organizational structure that is truly Byzantine. ICANN has so many "committees", "organizations", "working groups", "councils", and "assemblies" that one's mind goes numb simply looking at the fully detailed organizational chart. An older version of the org chart, one that lacks many of ICANN's newer elements, may be seen at http://www.icannwatch.org/images/orgchart.gif. Gilbert and Sullivan could easily write a sequel to The Mikado with ICANN as its subject.

The ever-ramifying complexity of ICANN's organization makes it exceedingly difficult for any but the most determined, or well financed, to penetrate past even the outer layers. This has made ICANN very much the province of professional business advocates and has deterred the participation of the average citizen.

It is frequently overlooked that ICANN in addition to its role over DNS, also regulates the allocation of Internet Protocol (IP) addresses. Address allocation policies will have a very significant impact on the future growth of the Internet and more particularly on what data carriers survive and dominate and which will fall by the wayside. It is likely that over the long term ICANN's IP address allocation policies will have a much greater economic impact than ICANN's Domain Name policies.

The very real technical need for IP addresses to be allocated and then sub-allocated in accord with the present day topology of the Internet creates a situation that tends to create a preferential lock-in for those who are currently at the top of the address allocation hierarchy and a discriminatory lock-out for those who may aspire to that role in the future. Address allocation is an area of substantial and subtle interaction between technical, economic, and social policies. In this area ICANN is for the moment leaving the task to the three worldwide regional address registries that were already doing the job before ICANN was formed.

How can ICANN obtain acceptance and legitimacy?

ICANN aspires to transcend national borders. ICANN conceives of itself as a supranational body that may act in ways that no single nation can, and equally, ICANN harbors a hope that it ought to be above the reach of the laws of any single nation.

And indeed it is true that ICANN has powers that supersede those of any single nation. For instance, ICANN's Uniform Dispute Resolution Policy (UDRP) amounts to a worldwide law, a law that is distinct and different from that enacted by any national legislature.

ICANN's UDRP is applied via a cascading contractual scheme. But because of ICANN's position as the sole gatekeeper of the Domain Name System, those who wish to have domain names have no choice but to submit to ICANN's will.

Under ICANN's UDRP trademarks are expansively interpreted at the expense of noncommercial uses of names and even traditionally acceptable nominative and free speech uses of trademarks. Suffice it to say, ICANN has created a new law of trademarks that as a practical matter overrides in many regards the trademark laws enacted by the Congress of the United States.

ICANN's decisions as to who does and who does not get a Top Level Domain (TLD) have transformed ICANN into an intrusive worldwide zoning board issuing licenses that determine who gets the privilege to set up a lucrative name-service shop on the Internet Boulevard.

Apart from the merits of the UDRP, this supranational scope is necessary - the Internet is supranational and ICANN's decisions as to its resources necessarily have supranational impact.

The real question is not whether ICANN ought to have this power but rather how ICANN's possession of that power obtains acceptance and legitimacy from the nations of the world and the users of the Internet.

My own answer is very simple: If ICANN

makes good decisions using sound procedures, it will come to be accepted as reasonable and legitimate.

As a Director I am very concerned that ICANN's rejection of public participation, its structural bias in favor of certain commercial interests, and its poorly defined and applied procedures will harm its long-term prospects for achieving such acceptance and legitimacy.

ICANN was given two distinct tasks by the US National Telecommunications and Information Administration (NTIA) when ICANN was formed. The first of these tasks was to deal with domain name issues. The second was to achieve public acceptance of this new thing under the sun. ICANN leapt into the first task - deciding DNS policy - while it slept on the second - achieving the public participatory structures that would provide a foundation upon which that DNS policy might be erected.

Until ICANN reforms its procedures and until it starts allowing meaningful public participation in its decision making, ICANN's policy decisions will be perceived as leaning towards special-interest concerns and thus undermine ICANN's long term hope of general acceptance and legitimacy. ICANN's obligation to have meaningful public participation in decision-making.

As mentioned above, ICANN's hopes for acceptance and legitimacy depend to a large extent on there being a perception that ICANN is responsive to all not just to some small set of business interests. To this end, a viable and believable means by which the public can participate in ICANN is essential.

ICANN is obligated to have a well-formed mechanism through which the public may meaningfully participate in ICANN's policymaking. However, ICANN has a history of impeding the creation of such a public role. Even today the public has obtained only a partial position - it elects only about one half of its nominated quota of Directors - and that partial position is at risk of being eroded or eliminated.

The "Executive Committee" of ICANN's Board of Directors appears to be increasingly active. The Executive Committee acts in lieu of the full Board, thus effectively eliminating any role for those Directors not on the Committee. The Executive Committee contains only one Director elected by the at-large membership. The impact of this is to dilute the role of the public by diluting the role of its elected representatives.

ICANN is explicitly required by its Articles of Incorporation to "operate for the benefit of the Internet community as a whole". [2] ICANN has made several promises to have a body in which the general public may fully participate in ICANN's policymaking. That promise remains unfulfilled.

In 1998 during ICANN's formative period, NTIA obligated ICANN to discuss various matters, including public participation in ICANN, with groups such as the Boston Working Group.

In 1999 ICANN formed a Membership Advisory Committee. This body issued its report in spring of 1999. The report was reasonably detailed and complete - and it favored the creation of a public "at-large" body that would elect Directors.

In July 1999, ICANN's chairman promised a subcommittee of the US House of Representatives that nine board seats would be filled by public election by the fall of the year 2000. [4] That promise was made to Congress 18 months ago. Although we did in fact have an election, it was for a mere five of the nine seats. Furthermore ICANN is now taking the position that the entire at-large body, including its quota of Directors, may be eliminated altogether.

Moreover, this election was not held until two years after ICANN's formation and well after many of ICANN's fundamental and important decisions had been made and put into effect.

The remaining four seats continue to be occupied by unelected people who were chosen, somehow, back in 1998 for one year terms. Those terms have now been extended to be at least four years. This extension is so long that my own term, and the term of all five of the Board members who were actually elected, will expire before then.

ICANN has initiated a "clean sheet" study to reconsider even the existence of the "at-large" membership and the election of board members by the public. I personally consider this study to be unnecessary as it does little more than revisit the ground already covered by ICANN nearly two years ago. I also consider this study to be overbroad because it explicitly places at risk even the bare existence of any public participation in the selection of ICANN's Board of Directors. This risk is not idle conjecture - within the last few weeks, ICANN executives have declared that the at-large membership no longer exists under ICANN's By-Laws.

Even if ICANN eventually implements a full quota of seats for the at-large, there will not be an election until fall of 2002 - four years after ICANN's inception. And in those four years, ICANN and time will have poured large amounts of metaphorical concrete around its prior decisions, making them essentially irreversible and forcing the public simply to accept that which was done while they were locked outside of ICANN's primary decision making body.

Changing the subject slightly - That election of the fall of 1999 had some strange characteristics and failures.

Many voters were denied the ability to register to vote, or if registered, were not given sufficient information and pass codes in order to cast their vote.

Neither candidates nor voters were allowed access by ICANN to the voter lists. This made it nearly impossible for the voters to discuss matters except via the limited channels provided by ICANN. This limitation made it virtually impossible for the voters to form coalitions or parties, to otherwise organize their votes, or to promote their favored candidates. In the long term, this will damage the ability of ICANN's voters to evolve into a well-structured and principled institution. Moreover, ICANN's denial of the voter lists was arguably in contravention to the Corporations Code of the State of California under which ICANN is incorporated.

I was elected by the at-large voters. But because of ICANN's restrictive controls over the voter rolls, ICANN has made it impossible for me to speak to my constituents or to solicit their advice.

This denial of the voter lists is justified by ICANN on the basis of privacy. Yet, the California legislature has determined that in corporate elections, the integrity of the election process requires that voters and candidates have means to communicate with one another outside of the view and potentially manipulative control of corporate management. If ICANN has a problem with the enactments of the California legislature, ICANN ought to take it up with the legislature rather than unilaterally undermining the viability of public participation in ICANN. ICANN's structure, whether taken piecemeal or as a whole, seems designed to include selected business interests - particularly those of trademark owners and DNS name registry/registrars - and to exclude Internet users. Deployment of a fully empowered at-large membership, with its full quota of Directors would go a long way towards redressing this imbalance.

Specific things Congress, the US Department of Commerce, and ICANN ought to do.

It is my desire to improve ICANN. To that end let me make some specific suggestions.

1. Congress should take care that the Internet does not serve as a means by which Federal administrative agencies slip their leash and assume unwarranted and undelegated powers.
2. Congress should take care that Federal administrative agencies do not try to do an end-run around their limited powers by outsourcing jobs to private bodies such as ICANN.
3. The Department of Commerce should exercise its independent judgment when dealing with recommendations coming from ICANN even if this may mean that the Department has to engage in hearings or other procedures.
4. The Department of Commerce should make it clear to ICANN that it expects ICANN to remember the obligations imposed on ICANN during its creation and thus improve its procedures and quickly create a fully formed vehicle for meaningful public participation in all of ICANN's decisions.
5. ICANN should be made accountable to someone more than just the Attorney General of the State of California.
6. ICANN should return to its mission and focus on technical coordination, leaving the public policy decisions to institutions better designed to accommodate public policy debate.
7. ICANN should fully adhere to the ideas of open access to all interested persons, transparent decision-making processes, and accountable decision makers. No ICANN process or body should be closed except when dealing with personnel, contract negotiations, litigation, or other expressly enumerated matters.
8. ICANN should emphasize implementation and deployment of good, fair procedures, such as its internal review mechanisms, even at the risk of delaying substantive policy decisions.
9. ICANN should follow the procedures written into its By-Laws and avoid ad hoc processes. In particular, this means more delegation of issues to the ICANN's specialized "supporting organizations."
10. ICANN should take steps to remedy the apparent capture by certain industry segments of ICANN's Domain Name Supporting Organization (DNSO).
11. ICANN should remove policymaking discretion from "staff" and sharply reduce the discretionary powers of executive officers.
12. ICANN should drop the "clean sheet" study of the at-large membership and simply get on with the job of filling all nine of the Board seats long promised to the public. At the same time ICANN should fully recognize the rights of at-large members as provided for under the California Corporations Code.
13. ICANN should rid itself of its excessively complex organizational structure.
14. ICANN should not have embedded entities that have no distinct legal status but which block review by members of the Board of Directors or the public.
15. ICANN should adopt better procedures for internal decision-making. In particular it should mandate semi-formalized procedures and rules of order for use by its numerous organizational entities.

Notes

[1] ICANN's purposes are stated in its paragraph 3 of its Articles of Incorporation. Emphasis has been added to highlight those parts mentioned in the text above.

3. This Corporation is a nonprofit public benefit corporation and is not organized for the private gain of any person. It is organized under the California Nonprofit Public Benefit Corporation Law for charitable and public purposes. The Corporation is organized, and will be operated, exclusively for charitable, educational, and scientific purposes within the meaning of §501 (c)(3) of the Internal Revenue Code of 1986, as amended (the "Code"), or the corresponding provision of any future United States tax code. Any reference in these Articles to the Code shall include the corresponding provisions of any further United States tax code. In furtherance of the foregoing purposes, and in recognition of the fact that the Internet is an international network of networks, owned by no single nation, individual or organization, the Corporation shall, except as limited by Article 5 hereof, pursue the charitable and public purposes of lessening the burdens of government and promoting the global public interest in the operational stability of the Internet by (i) coordinating the assignment of Internet technical parameters as needed to maintain universal connectivity on the Internet; (ii) performing and overseeing functions related to the coordination of the Internet Protocol ("IP") address space; (iii) performing and overseeing functions related to the coordination of the Internet domain name system ("DNS"), including the development of policies for determining the circumstances under which new top-level domains are added to the DNS root system; (iv) overseeing operation of the authoritative Internet DNS root server system; and (v) engaging in any other related lawful activity in furtherance of items (i) through (iv).

[2] The beneficiaries of ICANN's operations are described in paragraph 4 of ICANN's Articles of Incorporation. Emphasis has been added.

4. The Corporation shall operate for the benefit of the Internet community as a whole, carrying out its activities in conformity with relevant principles of international law and applicable international conventions and local law and, to the extent appropriate and consistent with these Articles and its Bylaws, through open and transparent processes that enable competition and open entry in Internet-related markets. To this effect, the Corporation shall cooperate as appropriate with relevant international organizations.

[3] ICANN's legal form is that of a corporation organized under the public-benefit/non-profit corporations laws of the State of California. ICANN is also exempt from US Federal taxes under 501(c)(3) of the tax codes.

[4] Esther Dyson, Chairman ICANN, made the following statements on July 22, 1999 before the House Committee on Commerce Subcommittee on Oversight and Investigations. (Emphasis has been added.)

Elected Board members. ICANN's elected Directors will join the Board in two waves: the first wave will consist of nine Directors chosen by ICANN's Supporting Organizations; the second wave will be elected by an At-Large membership consisting of individual Internet users. The Board expects the first wave to be completed by November 1999, and the second wave as soon as possible following that. In any event, the process of creating a fully elected Board must be completed by September 2000.

As to the first wave of elected Board members, ICANN expects that the nine Directors to be elected by its three Supporting Organizations (the Domain Name Supporting Organization, the Address Supporting Organization, and the Protocol Supporting Organization) will be selected and seated in time for ICANN's annual meeting in November in Los Angeles.

As to the second wave, it is ICANN's highest priority to complete the work necessary to implement a workable At-Large membership structure and to conduct elections for the nine At-Large Directors that must be chosen by the membership. ICANN has been working diligently to accomplish this objective as soon as possible. The Initial Board has received a comprehensive set of recommendations from ICANN's Membership Advisory Committee, and expects to begin the implementation process at its August meeting in Santiago. ICANN's goal is to replace each and every one of the current Initial Board members as soon as possible, consistent with creating a process that minimizes the risk of capture or election fraud, and that will lead to a truly representative Board.

ICANN Governance

Prepared Statement of A. Michael Froomkin Professor of Law University of Miami School of Law P.O. Box 248087 Coral Gables, FL 33124 before the Senate Commerce, Science and Transportation Committee Communications Subcommittee

Mr. Chairman and members of the Subcommittee, my name is Michael Froomkin. I would like to thank the Subcommittee for inviting me to appear today at this hearing on ICANN governance. I commend the Subcommittee for its wisdom and foresight in recognizing the importance of this issue.

I believe it is useful to separate this complex issue into three parts: (1) ICANN's mission or, if you will, ICANN's "jurisdiction"; (2) ICANN's internal organization; (3) The extent to which ICANN is subject to oversight by the Commerce Department, the U.S. Congress, or any other outside forces.

These three issues are intertwined. The nature and extent of ICANN's powers over the Internet and over Internet users that determines the type of internal governance structures which are appropriate for it. Similarly, the nature and quality of both ICANN's powers and its internal representativeness, not to mention checks and balances, determines the extent to which it needs to be subjected to searching external oversight. In particular, it is appropriate for this committee to enquire into the nature of the workings of the relationship between the Department of Commerce and ICANN.

Summary of Testimony

ICANN's go-very-slow policy on new gTLDs had no technical basis. Why then would ICANN adopt such a policy? The reason is that ICANN's policies are a product of an internal deliberative process that under-weighs the interests of the public at large and in so doing tends towards anti-competitive, or competitively weak, outcomes skewed by special interests

ICANN routinely claims to be either a technical standards body or a technical coordination body. If this were correct, then it might be proper for the Department of Commerce to defer to ICANN's presumed technical expertise and rely on ICANN's standards or allocation decisions without undertaking independent Administrative Procedure Act (APA) -compliant processes of its own. When, however, ICANN acts as policy-making rather than a standard-making body, then due to ICANN's unrepresentative nature its decisions do not carry any presumption of regularity or correctness and the US Government cannot rubber-stamp its decisions without additional independent fact-finding and deliberation.

We would all be better off if ICANN could confine itself to true standards issues, or to true technical coordination. If ICANN cannot, then ICANN needs to be subjected to constant scrutiny.

Terminological note: A "registrar" is a firm that contracts with clients ("registrants") to collect their information and payment in order to make a definitive and unique entry into a database containing all domain names registered in a top-level domain (TLD). This database is maintained by a "registry." Top-level domains are sometimes grouped into "generic TLDs" (gTLDs), which are currently three- or four-letter transnational domains, and "country code TLDs" (ccTLDs) which are currently two-letter TLDs. The "root" is the master file containing the authoritative list of which TLDs exist, and where to find the authoritative registries that have the data for those TLDs. Registrants typically register second-level domains (e.g. myname.com), but sometimes are limited to third-level domains (e.g. myname.genericword.com).

I. ICANN's Mission

ICANN's processes little resemble either standard-making or technical coordination. To date, ICANN's "standard making" has produced no standards. ICANN's "technical coordination" has been neither technical nor has it coordinated anything. Rather, in its initial foray into the creation of new gTLDs, ICANN has acted like a very badly organized administrative agency. Instead of engaging in standards work, ICANN is instead engaged in recapitulating the procedural early errors of federal administrative agencies such as the Federal Communications Commission (FCC).

What real standard-making would look like

A standard-based (or, at least, standardized) approach to gTLD creation would required ICANN to craft a pre-announced, open, neutral, and objective standard of competence rather than to pick and choose among the applicants on the basis of the ICANN Board's vague and inconsistent ideas of aesthetic merit, market appeal, capitalization, or experience. All applicants meeting that standard would be accepted, unless there were so many that the number threatened to destabilize the Internet (as noted below, if there is such a number, it is very large). ICANN might also put in reasonable limits on the number of TLDs per applicant, and on sequencing, in order to keep all of them going online the same day, week, or month.

Under a standards-based approach ICANN would have tried to answer these questions in the abstract, before trying to hold comparative hearings in which it attempted to decide to which of specific applicants it should allocate a new gTLD registry:

What is the minimum standard of competence (technical, financial, whatever) to be found qualified to run a registry for a given type of TLD? What open, neutral, and objective means should be used to decide among competing applicants when two or more would-be registries seek the same TLD string? What are the technical limits on the number of new TLDs that can reasonably be created in an orderly fashion per year? What open, neutral, and objective means should be used to decide among competing applicants, or to sequence applicants, if the number of applicants meeting the qualification threshold exceeds the number of gTLDs being created in a given year?

Today, reasonable people could no doubt disagree on the fine details of some of these questions, and perhaps on almost every aspect of others. Resolving these issues in the abstract would not necessarily be easy. It would, however, be valuable and appropriate work for an Internet standards body, and would greatly enhance competition in all the affected markets.

Once armed with a set of standards and definitions, ICANN or any other allocation body, would be on strong ground to reject technically incompetent or otherwise abusive applications for new gTLDs, such as those seeking an unreasonably large number of TLDs. A thoughtful answer would inevitably resolve a number of difficult questions, not least the terms on which a marriage might be made between the Department of Commerce's "legacy" root and the so-called "alternate" roots.

What technical coordination would look like

An alternate approach to gTLD creation, one that would most certainly enhance competition, would take its inspiration from the fundamental design of the Internet itself-and from major league sports. The Internet was designed to continue to function even if large parts of the network sustained damage. Internet network design avoids, whenever possible, the creation of single points of failure. When it comes to policy, however, ICANN is currently a single point of failure for the network. A solution to this problem would be to share out part of ICANN's current functions to a variety of institutions.

In this scenario, ICANN would become a true technical coordination body, coordinating the activities of a large number of gTLD policy partners. ICANN's functions would be: (1) to keep a master list of TLDs, (2) to ensure that there were no 'name collisions' - two registries attempting to mange the same TLD string; (3) to fix an annual quota of new gTLDs; (4) to run an annual gTLD draft; (5) to coordinate the gTLD creation process so that new gTLDs came on stream in an orderly fashion instead of all at once.

Each of ICANN's policy partners would be assigned one or more draft choices, and then ICANN would randomly (or, perhaps, otherwise) assign each one their draft picks. As each policy partner's turn came up, it would be entitled to select a registry - imposing whatever conditions it wished - to manage any gTLD that had not yet been claimed on ICANN's master list. In keeping with the transnational and public/private nature of the Internet, ICANN's policy partners could be a highly diverse mix of international, national, and private "civil society" bodies.

While I think this alternate solution would best achieve the ends of internationalization, competition, and diversity, it might well require legislation since it is unclear if the Department of Commerce has the will (or the authority) to implement such a plan, and we have seen no sign that ICANN is about to divest itself of any policy authority unless forced to do so.

What ICANN actually did: select an arbitrarily small number of gTLDs based on arbitrary appraisals of aesthetic merit, market appeal, capitalization, and experience.

Rather than adopt either a standards or a technical coordination approach, ICANN instead adopted an arbitrary approach. First it set an arbitrarily low ceiling on the number of TLDs, then it allocated most but not all of that quota based on its arbitrary appraisals of the applicants aesthetic merit, market appeal, capitalization, and experience.

ICANN's decision to impose an arbitrary limit on the number of new gTLDs

The closest thing to technical standards work that ICANN has done to date was to adopt an artificially low limit on the number of gTLDs it would recommend the Commerce Department create - under the guise of a so-called "proof of concept". The grounds on which ICANN based this arbitrarily low limit on the number of new gTLDs demonstrate as clearly as anything else that ICANN is not making technical decisions but instead making policy choices on the basis a wholly inadequate an unrepresentative structure.

ICANN has never claimed that the technical stability of the DNS would in any way be threatened by the introduction of a very large number of new gTLDs. Indeed, it could not easily make this claim, since all the technical evidence is to the contrary. Rather, the dangers that ICANN seems concerned about are social - potential consumer confusion, and a potential 'land rush' mentality due to the enormous pent-up demand. (In my opinion, however, ICANN has selected a policy that maximizes the risk of a 'land rush'. Panic buying happens when consumers fear a shortage. Here, ICANN is proposing the creation of a very small number of gTLDs, with no assurances as to when if ever the next batch will be created. This gets it exactly backwards: the way to avoid a land rush would be to have a very predictable path for new gTLDs so that everyone understands that there's no need to panic since plenty of names will always be available.)

I am not an expert on Internet engineering. However, my understanding is that while experts do not agree on precisely how many gTLDs could be created without adverse consequences to DNS response time, there appears to be a technical consensus that we are nowhere near even the lowest possible limit. ICANN At-Large Director Karl Auerbach, himself a technical expert, has suggested that the smallest technically-mandated upper level for the number of gTLDs might be as high as a million. (1) Persons with long experience in DNS matters, including BIND author Paul Vixie, apparently agree. (2) Others have performed tests loading the entire .com file as if it were a root file, and found that it works. In principle, this is not surprising, as there is no technical difference between the root file containing the information about TLDs and a second-level domain file. Given that there are currently about sixteen million registrations in .com, if this argument is right, then the maximum number of TLDs may be very high. (3) Some experts worry, however, that a very large number of new TLDs, such as a million, might affect DNS response time. (4) If so, that still means that with fewer than 300 TLDs in operation today (gTLDs + ccTLDs), we can afford to create tens of thousands, and probably hundreds of thousands, more.

It is an article of faith among Internet entrepreneurs that possession of a good domain name is a necessity for an Internet startup. Many traditional firms also consider the acquisition of a memorable or short domain name to be of strategic importance. Recently, for Internet startups, possession of a "good" name was seen as a major asset - reputedly enough in some cases to secure venture financing.

For some time now, however, it has also been an article of faith in the Internet community that "all the good names are taken" Recently it has seemed as if simply all the names that were a single word were taken. This apparent shortage, especially in .com, has driven firms seeking catchy names into the aftermarket. There does appear to be a reasonably large stock of names in the existing gTLDs being held by domain name brokers for resale in the aftermarket. Prices are very variable. Although few firms paid millions of dollars like the purchasers of business.com, and loans.com, it appears that at least until the .com bubble burst, the shortage of attractive names in .com , and the resulting need to purchase them at high markups in the aftermarket created what amounted to a substantial "startup tax" on new businesses.

ICANN justifies its very tentative initial foray into gTLD creation as a "proof of concept" but it has not disclosed the concept that is believes it is trying to prove, nor described how one tells if the test is successful, nor even when one might expect ICANN to do the evaluation. The "concept" cannot be gTLD creation itself: There is no rocket science to the mechanics of creating a new gTLD. From a technical perspective, creating a new gTLD is exactly like creating a new ccTLD, and creating new ccTLDs is quite routine. Indeed, .ps, a TLD for Palestine, was created less than a year ago with no noticeable effect on the Internet at all. (5)

In fairness, ICANN is not originally responsible for the gridlock in gTLD creation policy, which in fact long predates it. Indeed the Department of Commerce - which currently has the power to create new gTLDs -

called ICANN into being because it wanted to find a politically feasible way to create new TLDs in the face of difficult political obstacles, not least a belief in the intellectual property rights holders community that new TLDs might add to the risk of customer confusion and trademark dilution.

This political fear, more than any mythical technical consideration requiring a "test" or "proof of concept", explains why ICANN imposed a needlessly low limit on the number of new gTLDs it would recommend the Department of Commerce create in this first round, and why ICANN has as yet not been able to consider when if ever it will contemplate future rounds of gTLD recommendations. It does not explain, however, why ICANN why ICANN persists in falsely claiming consensus for its artificially low number of TLDs, nor why went about selecting its seven finalists in the manner it did. Indeed, as described below, ICANN's gTLD selection procedures were characterized by substantial failures.

Nevertheless, it might seem that despite any procedural irregularities, ICANN's recommendation that the Department of Commerce create a small number of new gTLDs can only be good for competition as it will increase supply and thus drive down prices. And indeed, supply will increase. Unfortunately, of the new gTLDs, only .biz and maybe .info are likely to be of attractive to the majority of startups and other Internet newcomers. Because there are only two such domains, and because there is no easily foreseeable date at which additional gTLDs might become available, there is a substantial risk of a speculative frenzy in which domain name brokers, cybersquatters, and amateur arbitragers all seek to register the catchy names that have not already been snapped up by trademark holders who took advantage of their pre-registration period. I am concerned that the faction which controls ICANN will use this very predictable speculative frenzy as 'evidence' that new gTLDs are a bad idea, or that the number must be kept down in the future.

The surest way to drive down and keep down the price of domain names, thus eliminating the "startup tax" and enhancing the ability of new firms to enter new markets and incidentally greatly reducing, perhaps even almost eliminating, cybersquatting, is to create healthy expectations. As soon as participants in the market understand that a steady supply of new domain names in attractive gTLDs will continue to become available on a predictable schedule, the bottom will fall out of the after-market, and the incentive (albeit not the opportunities) for cybersquatting will be greatly reduced, thus helping e-commerce by making attractive names available on reasonable terms to a much greater number, and wider variety, of persons and firms.

Selection of gTLDs

In ICANN's recent gTLD process, ICANN acted not as a standards or coordination body, but as if it were allocating scarce broadcast spectrum is some kind of comparative hearing process. ICANN created no standard. It 'coordinated' no projects with running code being deployed by outside parties. Rather, ICANN acted like a foundation grant committee, trying to pick 'winners.' In practice, ICANN's exercise of its gatekeeper committee role contributes to the artificial shortage of gTLDs. Worse, the selection processes ICANN employed were amateurish and arbitrary.

Although all applicants were charged the same non-refundable $50,000 fee, a sum that immediately skewed the process towards commercial uses and away from non-profit or experimental uses, it appears not all applicants received equal treatment. During the Los Angeles ICANN Board Meeting, it transpired that the staff had not subjected all the proposals to the same level of analysis. Thus, when Board members sought more detailed information about proposals that interested them, but which the staff had relegated to the second tier, that information sometimes did not exist, although it existed for the staff's preferred picks.

ICANN then attempted to hold a one-day comparative hearing between more than 40 applicants, each of whom had complex applications that referenced multiple possible gTLDs. During this process, each applicant was given three minutes to speak.

Both before and during the one-day Board meeting, both the staff and the Board seemed excessively concerned with avoiding risk. Although true competition in a fully competitive market requires that participants be allowed to fail if they deserve to do so, there are reasonable arguments as to why it makes sense to have a body like ICANN require potential registry operators to meet some minimum standard of technical competence. One can even make a case for requiring a showing of some financial resources, and for requiring the advance preparation of basic registry policy documents spelling out who will be allowed to register names and under what terms. Perhaps there are other neutral criteria that should also be required and assessed. This is a far cry from ICANN's apparent tendency to tend to prefer established institutions and big corporations, and to downplay the value of experience in running code. If in 1985 the Internet itself had been a proposal placed before a committee that behaved as ICANN did in 2000, the Internet would have been rejected as too risky. Risk aversion of this type is antithetical to entrepreneurship and competition.

Worst of all, ICANN applied its criteria arbitrarily, even making them up as it went along. The striking arbitrariness of the ICANN decision-making process is illustrated by the rejection of the ".union" proposal based on unfounded last-minute speculation by an ICANN board member that the international labor organizations proposing the gTLD were somehow undemocratic. (That this same Board member was at the time recused from the process only adds to the strangeness.) The procedures ICANN designed gave the applicants no opportunity to reply to unfounded accusations. ICANN then rejected ".iii" because someone on the Board was concerned that the name was difficult to pronounce, even though the ability to pronounce a proposed gTLD had never before been mentioned as a decision criterion. I am not in a position to vouch for the accuracy of each of the claims of error made by the firms that filed reconsideration requests after the Los Angeles meeting (available at http://www.icann.org/committees/reconsideration/index.html) but as a group these make for very sobering reading.

If ICANN were to limit itself to either standard making or technical coordination it would have approached its mission very differently from the arbitrary and amateurish procedures it used. It is critical to note that the relevant standards of comparison for ICANN's decision making are not the private sector. As a non-profit standards body contracting with the US government, ICANN should either be held to standards of openness, professionalism, and neutrality appropriate for standard-making or, if making political and social choices, be treated as a state actor and expected to act in conformity with fundamental norms of due process. Suggestions heard from some victorious gTLD applicants that ICANN's processes compare favorably with those used for procurement in the private sector are both erroneous and irrelevant. ICANN is not engaged in procurement. It is not "buying" anything. And ICANN paid almost no attention to the prices proposed by would-be registries.

II. Internal Organization

ICANN's go-very-slow policy on new gTLDs had no technical basis. Why then would ICANN adopt such a policy? The

reason is that it is a product of an internal deliberative process that under-weighs the interests of the public at large and in so doing tends towards anti-competitive, or competitively weak, outcomes skewed by special interests.

The source of this predisposition is the distribution of decision-making authority on the ICANN Board, and in ICANN's subsidiary institutions, which have been manipulated to neuter the public voice, and the role of individuals, non-profits, and civil society groups. Originally, half of ICANN's governing Board would have been elected by at-large members of ICANN. Instead, ICANN has worked at every turn to prevent this.

In July, 1999, ICANN Chair Esther Dyson told the House Commerce Committee's Subcommittee on Oversight and Investigation that ICANN's "highest priority" was to elect nine at-large Board members, (6) exactly as ICANN had committed to do as an original condition of being approved by the Department of Commerce. Instead, ICANN reneged on its commitment to the United States government, and to the public, that half its Board would be elected by an at-large membership. Thus, today:

Instead of half (nine) of the Board members being elected at large, as promised to NTIA and to Congress, ICANN amended its by-laws to allow only five members to be elected at large; Instead of all the self-appointed nine original directors leaving office as they promised Congress and the public they would do, four remain in office; Instead of allowing the five elected at large members to participate in the selection of the new gTLDs, ICANN amended its by-laws to seat them at the close of a meeting, instead of at the start (the process used for all previous new directors). Then ICANN rushed its processes so that it could make the final decisions minutes before the new directors took office. In a move that risks further neutering the five elected at-large members, ICANN announced that their jobs would all be abolished at the end of their two-year terms, unless a majority of the full Board voted (after a "clean sheet study") to re-establish elected at-large Board seats.

[Note that under the current by-laws, the un-elected directors apparently get to keep their jobs indefinitely.] The internal institutions that ICANN created to take the lead in domain name policy - the seven constituencies in the "Domain Name Supporting Organization" (DNSO) - were designed from the start to exclude individuals from membership. The very engineers who built the Internet are not represented in their personal capacities - only if their employers choose to send them. All non-commercial groups, including all universities, all consumer groups, all political groups throughout the world are shoehorned a single DNSO constituency. They are, in the main, ineligible for full voting membership of any of the other six constituencies.; Meanwhile, many businesses such as Internet first-movers and others who have an interest in reducing online competition for established firms are eligible to be in two, three, or even four of the seven constituencies, thus allowing them multiple votes-and a certain majority.

The interest groups that acquired a voting majority in those institutions have shown relatively little interest in the rights and needs of small businesses, non-commercial entities, or individuals. They have shown considerably more interest in securing special protections for trademarks, above and beyond what is provided by statute, than they have in maximizing the liberty-enhancing and competitive potential of the Internet.

ICANN is a highly complex organization (see attached charts, prepared by Tony Rutkowski). It is simply impossible for anyone to keep track of what is happening in all the different pieces, except an organization capable of deploying a fleet of lawyers. Similarly, because ICANN sees its mission as global, it meets four times a year on four different continents. Next month's meeting, for example, is in Australia. The result of this laudable attempt at internationalization is that only interests wealthy enough to attend all these meetings - with several representatives - can achieve the continuity of participation required to influence ICANN's decisions in any sort of a consistent manner. The result tends to be a 'consensus' of those with the necessary expense accounts.

III. External Checks on ICANN

I do not deny that one can identify potentially serious social issues that might be caused as side effects of the creation of new gTLDs. I do submit that ICANN has no competence to deal with them, and that its actions have to date in creating special domain name registration rights for trademark holders, well in excess of the rights granted to them by Congress, have been anti-competitive, unfair, and counterproductive.

ICANN's mandate and its competence is, at most, for technical matters. Social policy issues such as the intellectual property consequences of new gTLDs, the number of days a person should have to respond to an arbitration over a domain name, or issues of content management, should not be decided by engineers or by the people who happen to have seized control of ICANN. Rather, they should be decided via the means we traditionally use for making social policy choices - markets and representative democracy.

Since ICANN's decisions as to its gTLD recommendations were not based on purely technical criteria, as a formal matter ICANN is making social policy choices, not just acting as a standards body. It is therefore right that ICANN's decisions are subject to external checks. Indeed, as I argue in my article Wrong Turn in Cyberspace: Using ICANN to Route Around the APA and the Constitution, 50 Duke L.J. 17 (2000), available onlinehttp://www.law.miami.edu/~froomkin/articles/icann.pdf , as a matter of law ICANN as currently constituted amounts to a state actor, and thus is subject to the same Due Process constraints as apply to any federal agency. Accordingly, its arbitrary and capricious decisions violate both the APA and the Due Process Clause of the Constitution.

ICANN and the U.S. Department of Commerce dispute this characterization. They prefer to rely on form over reality, and insist that ICANN is legally private despite the fact that ICANN derives all of its authority and revenue from Commerce's loan to ICANN of authority over the root. It follows, however, that if this characterization of ICANN as a purely private body is correct, then there are strict limits on the extent to which the Department of Commerce can implement ICANN's recommendations without violating the Administrative Procedures Act, or the Constitution's Due Process clause.

Once ICANN makes its formal recommendations, the Department of Commerce will have to decide how to proceed. Rubberstamping of ICANN's decisions by the Department of Commerce would amount to adopting ICANN's arbitrary and capricious choices, since the U.S. government would essentially endorse both ICANN's practices and its conclusions.

The Department of Commerce has maintained that its relations with ICANN are not subject to the APA, or indeed to any legal constraint other than those relating to relations with a government contractor and/or a participant in a cooperative research agreement. But whatever the legal arguments, when contemplating decisions which will shape the very nature of the Internet naming system, Commerce should proceed with deliberation, and act only on the basis of reliable information. The need for reliable information, proper public participation, and transparent and accountable decision-mak-

ing is even stronger when Commerce contemplates making the sort of social policy choices - as opposed to mere technical standard-setting - embodied in creating new gTLDs and imposing conditions on their use. Basic requirements of fairness, due process, and the need to make reasonable decisions counsel in favor of notice, public access, the making of an official record, and deliberation.

There is no question but that if a federal agency had acted as the ICANN Board did, its decisions would not satisfy even cursory judicial review. In the circumstances, therefore, it would be unreasonable and a denial of due process for Commerce to rely on the outcome of such a flawed process without conducting its own review.

ICANN faces a choice: On one path it becomes a true standards body, or a true technical coordination body, and leaves the social policy choices to those - like Congress - who have the legitimacy to make them. On the other path, the one it currently seems to be following, it is a state actor. In that case, its actions to date have been far too arbitrary to survive judicial review.

NOTES

1. Posting of Karl Auerbach, karl@CaveBear.com, http://www.dnso.org/wgroups/wg-c/Arc01/msg00195.html .

2. E-mail from Paul Vixie, BIND 8 Primary Author, to Eric Brunner (Dec. 15, 1999) ("A million names under '.' isn't fundamentally harder to write code or operate computers for than are a million names under 'COM.'"), http://www.dnso.org/wgroups/wg-c/Arc01/msg00203.html .

3. See Quickstats, at http://www.dotcom.com/facts/quickstats.html (reporting twenty million registrations, of which 80% are in .com).

4. See, e.g., E-mail from Paul V. Mockapetris, BIND Author, to Paul Vixie, BIND 8 Primary Author, & Eric Brunner (Dec. 15, 1999) (querying whether one million new TLDs would impose performance costs on DNS), http://www.dnso.org/wgroups/wg-c/Arc01/msg00202.html .

5. See IANA Report on Request for Delegation of the .ps Top-Level Domain, at http://www.icann.org/general/ps-report-22mar00.htm (Mar. 22, 2000).

6. Testimony of Esther Dyson, Chair, ICANN, before the House Commerce Committee, Subcommittee on Oversight and Investigations, July 22, 1999, http://www.icann.org/dyson-testimony-22july99.htm .

Testimony of the Domain Name Rights Coalition and Computer Professionals for Social Responsibility

Introduction

Thanks to the Committee for providing the opportunity to provide feedback to the Senate regarding the role of ICANN and the Commerce Department in the ongoing battle for Internet governance. Although you have received letters from others who attempt to downplay ICANN's role, make no mistake; it goes far beyond that of technical management and enters the realm of a regulatory body. ICANN's policy will affect commerce, freedom of expression, and likely stifle the very medium it seeks to regulate. ICANN has not provided an accurate picture of the Internet world to the Committee. We felt it was necessary to correct and explain much of what they reported to you in response to your questions.

About DNRC and CPSR

The Domain Name Rights Coalition has participated in the ongoing debates concerning Internet management as a member of the Boston Working Group, a member of the Open Root Server Confederation, former steering committee member of the IFWP (International Forum on the White Paper.) DNRC submitted comments on the Green Paper, use of the .US domain, testified before Congress, submitted comments to the World Intellectual Property Organization, and has dissented in the formation of ICANN's Uniform Dispute Resolution Policy.

CPSR is a public-interest alliance of computer scientists and others concerned about the impact of computer technology on society. We work to influence decisions regarding the development and use of computers because those decisions have far-reaching consequences and reflect our basic values and priorities.

As technical experts, CPSR members provide the public and policymakers with realistic assessments of the power, promise, and limitations of computer technology. As concerned citizens, we direct public attention to critical choices concerning the applications of computing and how those choices affect society.

Summary

ICANN continues to execute fundamental Internet policies beyond its mandate as "technical coordinator," and without creating the participatory structures that would allow its decisions to be accepted and trusted by a broad spectrum of stakeholders.

The sad fact is that ICANN has been "captured" from the beginning. Special interest groups have dictated the direction of ICANN, and have morphed it into an Internet Governance body with none of the protections afforded by governments.

Governmental safeguards to American ideals such as Free Speech and other civil liberties, must be codified in ICANN, as well as other quasi governmental corporations in the private sector.

There is no technical reason to refuse any applicant for a top level domain. Instead, policy reasons were substituted for technical reasons, resulting in limiting competition, not enhancing it.

By maintaining a false artificial scarcity, ICANN is risking an increasingly fragmented and incoherent Internet system. By their own statements ICANN claims to be concerned with stability above all. However, ICANN has now actively sought to cause domain names already registered by existing businesses (.web, .biz, .museum, .pro, .info and others) to be registered to potentially different parties at another. Rather than taking the opportunity to strengthen the domain name system, ICANN is risking the single predictable factor of the Internet. They are, in effect, ensuring that current domain names maintain their scarcity, and thus value. The Department of Commerce hopefully did not contemplate that ICANN would become the Federal Reserve Board of domain names.

Domain name registrars who have proven technical competence beyond a doubt by registering domain names for years as well as providing their own alternative roots were refused permission, not on technical grounds, but purely on policy grounds. TLD registries should be allowed to set policy independently of ICANN and that no registry be excluded from TLD operator status if

its policy differs from that of ICANN.

The burden of proof should be placed on ICANN to refuse to admit competition. Potential competitors should not prevented from entering the market, and force to prove to ICANN that they are worthy of an artificiality small number of slots.

ICANN further claims to honor intellectual property law. Yet the fundamental basis of a natural right of property is that one earns property by the sweat of the brow. The ICANN uniform dispute resolution protocol (called a protocol to falsely deny that it is clearly a policy document, with negligible technical content) does not acknowledge any sweat of the brow argument. In multiple cases (for example, etoys and workingwomen) entrepreneurs entered the risky world on on-line commerce and sweat over their domain names to create value. The UDRP policy does not acknowledge that risk-taking or investment even over such generic words and phrases as "toys" and "working women." The UDRP appears to have nothing to do with law. For example the published procedural rules used by one registered UDRP provider are in clear violation of any standard of procedural due process.

As currently constituted ICANN has failed on all charges. It has moved slowly; been unrepresentative; acted to limit competition; and failed to offer useful, fair, coherent policies, or even policies which encourage investment in virtual property. ICANN is a policy experiment that has failed.

History

The Domain Name Rights Coalition was formed in 1995 directly because of the NSI domain name dispute policy which we thought stifled the rights of individuals and small businesses to choose domain names. The development and growth of the World Wide Web brought with it a significant interest by the business community. It soon became clear that IANA, a US government contractor run by Dr. Jon Postel, would be unable to continue its management of domain names and numbers without significant help. The first attempt to transfer control occurred in 1994 when Dr. Postel attempted to place IANA under the Internet Society (ISOC.) This failed, but something else grew from that union. The IAHC (International Ad Hoc Committee) was created, and tried to take over Internet governance via a document called the gTLD-MOU. Comments were solicited by the IAHC from the Internet community, but the responses were largely ignored. It is not coincidental that many of the members of CORE, POC (the Policy Oversight Committee) , ISOC (an original IAHC advocate), WIPO, and the ITU are now heavily involved with the ICANN process, and have in a sense "captured" that process.

The gTLD-MOU was stopped by the Internet community when it became clear that the process was closed, unaccountable, and non-transparent. Various people appealed to the Department of Commerce and the State Department for help. Through significant work and effort, the IAHC plans were thwarted, and the Commerce Department produced the "Green Paper" as a roadmap for technical management of names and numbers. The Green Paper was truly a pro-competitive solution, one that was hotly contested by many European Governments, and the previous supporters of the MoU. In fact, it was right around this time, that Jon Postel redirected over half of the world-wide root servers to his server in California. While we may never know, this combination of events apparently derailed the Green Paper, and started the process that resulted in the White Paper.

Thousands of comments were submitted by a large cross section of the Internet community, although many questioned (and still question) under what authority the Department of Commerce was taking control of Internet functions. Many of these comments were incorporated in the "White Paper" which provided a framework for considering these issues. Using the White Paper as a foundation, the IFWP (International Forum on the White Paper) was created in 1998 to discuss these issues and attempt to reach the consensus that was required to move forward with the plans envisioned in the White Paper for an open, transparent and accountable organization, Newco, to manage domain names and numbers. Please note that even with the White Paper, significant numbers of people still ask under what authority Commerce is operating in choosing one company over another, mandating that company's bylaws, mandating that company to be non-profit, and assisting in choosing the unelected board members of that company.

The IFWP steering committee consisted of members of the Internet community who were involved with not-for-profit enterprises. These included CORE, the Commercial Internet Exchange (CIX), Educause, the Domain Name Rights Coalition (DNRC), and various other groups. It was chaired by Tamar Frankel, a respected law professor and expert on corporate structure and process from Boston University. The IFWP held meetings around the world, and worked to come to consensus on various issues. In the midst of this process, Joe Sims, attorney for Dr. Postel, prommulgated a set of by-laws for Newco. He did this in closed meetings with no public input. These by-laws were presented to the IFWP, but did not gain consensus, largely because the points on which the IFWP had already garnered agreement were not included. Various further drafts followed, but still none of them achieved consensus.

In late August of 1998 after the final IFWP meetings, the steering committee met telephonically to plan the final or "wrap up" meeting in which the consensus points would be memorialized, and further concessions would be provided by all sides. Although there had been multiple votes already taken that clearly supported a wrap up meeting, yet another vote was called at that time. Mike Roberts vehemently opposed a wrap up meeting, and was supported in this by Barbara Dooley of the CIX. There is speculation that Mr. Roberts had already been contacted at that time regarding serving with the ICANN board in some capacity. Further, around the time of the wrap up meeting, Esther Dyson says that she was approached by Roger Cochetti of IBM and Ira Magaziner in Aspen, Colorado and asked if she would be interested in joining the ICANN Board. The IFWP wrap up was finally completely derailed by ICANN's refusal to participate in the meeting.

Some of the members of IFWP continued their work to create an open, transparent and accountable Newco. Two major groups, the Open Root Server Confederation (ORSC) and the Boston Working Group (BWG) promulgated by laws for Newco through open process. Three sets of by-laws were provided in a timely manner to the Department of Commerce. Although the Commerce Department had long stated that they would not choose one set of by-laws over any other, they chose the ICANN's bylaws as a starting point

The Commerce Department directed ICANN to consult with the BWG and the ORSC regarding areas of concern to Commerce but there was little reason for them to do so since their bylaws and structure had already been chosen. ICANN did meet telephonically with BWG and ORSC, but failed to make substantive changes in its bylaws to accommodate the diversity of opinions towards fundamental issues such as openness of board meetings, voting on the record, voices for individuals and non-commercial entities, limitations on ICANN's powers to strictly technical issues, etc. Both BWG and ORSC warned that the concept of constitu-

encies would lead to capture by corporate interests at the expense of expression. BWG wanted to do away with constituencies altogether. ORSC wanted constituencies structured so that everyone would have a voice. The ICANN constituency structure has, as predicted, become the catalyst for capture by the old gTLD-MOU crowd, and a large and powerful group of trademark interests. These trademark interests are currently pressing non legislative expansion of rights for trademark holders, at the expense of free speech and expression.

Competition

It is ironic in that in the midst of all the controversy over competition, ICANN has hesitated to take the single step that would introduce the most competition: creating objective technical guidelines for choosing new TLD registries. Although ICANN has indeed chosen 7 new gTLDs, no guidelines have been established that would allow for future expansion. There are no roadmaps by which prospective registries can turn to structure their technical business plans.

The Process of Consensus Development and Implementation

ICANN is correct in that its formation was an unprecedented experiment in private sector consensus decision-making. Unfortunately, that experiment is in the process of failure. ICANN's claim of "openness and transparency, based on Internet community consensus, bottom-up in its orientation and globally representative" is far from the reality of the situation. ICANN is the classic top-down organizational structure without accountability. When its by-laws are inconvenient, they are changed without discussion.

Board of Directors

Currently, the 9 seats that were to be elected from the Internet stakeholders, the so called "at large" directors, were whittled down to 5. The other 4 seats have been held by "Board Squatters," those who were appointed and not elected. Despite calls for elections to replace the squatters, and calls for their resignations, no movement has occurred.

Instead, the Board has responded with a "clean sheet" study that could, conceivably, dismantle the entire At Large process altogether. Leaving ICANN controlled solely by special interest groups.

ICANN Staff

ICANN's staff seems, by all outside examination, to be driving all policy decisions. The non elected staff, submits reports to the board which are normally accepted verbatim, with no indication to the Internet community of what criteria was used to reach the conclusions contained therein. These policy decisions, often clearly outside the reach of a "technical management" organization, are then presented as a "fait accompli" with no accountability or transparency, and no input from the Internet community that they affect.

Conclusion

The Internet is the single most significant communications medium ever created. Its power goes well beyond that of shopping malls and e-commerce, and empowers individuals in a way never before imagined. It is thus a national as well as an international resource. The ability to control important aspects of this technology cannot be underestimated. It is up to all of us to remain vigilant when organizations are given special privilege by a branch of the US Government to control this vast means of expression. Safeguards must be put into place whereby individuals, non-profit entities, churches, tribal governments, and other disenfranchised groups may provide unencumbered input and opinion to an open, transparent and accountable entity. This entity is, unfortunately, not ICANN in its current form.

ICANN must be restructured. We suggest the following changes:

1. ICANN must limit itself to technical coordination only. This limitation must be irrevocably codified in ICANN's bylaws, and must be enforced by the Commerce Department and/or Congress

2. All policy decisions, including the new selection of gTLDs must be clearly documented as to what objective criteria was used to select them. Any decisions without such objective, clearly stated criteria should be rescinded and revisited after such objective criteria are put in place.

3. ICANN's current constituency structure must be restructured to allow for more inclusion by Internet stakeholders, including individuals, educational entities, religious entities, consumer protection groups, civil libertarians, and others. The current practice of lumping all of these groups into one constituency, while leaving 6 others who all represent overlapping business interests, must change.

4. ICANN must not be used as the arm of government to circumvent constitutional rights and liberties. An example is the "takings" clause. Several gTLDs are being operated currently that will essentially be "taken" if ICANN puts the identical strings in their root system. Another example is ICANN's non accountability under the Federal Administrative Procedures Act.

5. ICANN must not be allowed to pick and choose provisions of its mandate that it will accept and others that it will ignore. The most glaring example is its lack of codifying the At Large group into an irrevocable part of the ByLaws. Second to this is ICANN's failure to recognize a place for individuals to participate on an equal footing with business interests. Third, is ICANN's continued failure to constitute a membership in accordance with the White Paper, as well as California public policy under which it is organized.

6. Fundamental rights of American Citizens, such as Free Speech must trump intellectual property rights of businesses. ICANN's Uniform Dispute Resolution Policy gives trademark and intellectual property holders a means to limit and silence legitimate speech without recourse. If ICANN is allowed to continue to use this policy, a balance must be struck whereby speech rights are protected and abuses by intellectual property holders are curtailed.

Mikki Barry

Continued from page 26

Valentines day. He put himself in the room only reluctantly out of determination that the organization wielded by Roberts iron hand would not destroy the tradition of Internet leadership by honest and open debate conducted on a level playing field by people who were motivated by the importance of building something bigger than they were and not by a crass quest for power and money .

The outcome of debate on St Valentines day between these two men will be very important for those who care about the preservation of a decentralized and open internet. One man who wants an open membership and the other who fears it.

Executive Summary

The Netherlands and SURFnet, pp. 1- 8

We interview Kees Neggers, Managing Director of SURFnet who explains first and foremost the government's policy that research and education community must pay for its own internet connectivity while government subsidies are used to "tender" for a commercially advanced network. What we did is to ask for things that we think are technically possible, but for whatever reason the commercial market is not yet willing to deliver."

He states that 'the nature of the resulting contract is more partnership for a common development than just a normal supply contract. Research people from both Telfort(BT) and Cisco Systems are actively involved in our activities." Since the late 1980s he has operated on the following dynamic: " with the 'innovation' money we built a network — SURFnet1. With the user money, we made it operational and kept it operational. In parallel, we used the government money to build the next generation network. Now we are building SURFnet5 with government money and the users are paying for SURFnet4 which is essential, because it is the lifeline to the Internet for our customers."

Kees' philosophy is that leading edge network infrastructure will enable many new and productive applications. How the bandwidth is produced is of primary importance as government money is used to catalyze efforts by the commercial carriers in binging new technologies to market.

SURFnet is asking for and receiving a similar kind of optical customer owned lambda network as Canarie in Canada. "In our next generation networking, we are no longer just talking to carriers to make plans, we are talking to customers to make plans. And the carriers are welcome to deliver what we ask. But you see, we have the strengths of the end user community to tell them what we want. As with SURFnet5, we told them what we wanted : 10 gig lambdas, rather than ATM and SDH."

The fiber builds that it undertakes for its backbone are also used to attract and create commercially created fiber infrastructure. SURFnet's philosophy has been one of continually pushing the envelope. To do this it needs to be able to move quickly and make its own decisions promptly. In part for these reasons it has chosen not to participate in the next generation pan European GEANT project which is a consortium of European national research and education networks.

"SURFnet has found that European networking so far tends to be driven by the speed of the slowest. I don't like that model. It is a recipe for not being able to be state of the art. We wanted to use the opportunity of the setting up of Geant to learn from the past and improve the situation. However the way the Geant network is organized is a continuation of the structure from the past. Dante has no central management. It's a consortium of some twenty-six national research networking entities. All 26 have to agree on everything. I didn't want to be the 27th of that group. So my proposal was, Dante should do the procurement, should operate the network and I want to be a customer of DANTE. And the consortium should only be a consortium in its relationship to the European Commission to define the set-up. But none of the participants in Europe were willing to go that route." While it is the only European nation not to participate in Geant, it is connected to the pan European Ten-155 research backbone and will connect to Geant as well.

SURFnet has been following the development of OBGP. According to Kees: "The reason that Bill St. Arnaud and I are interested in this is that we are not providers. We work for a user community. If it is better for the user community to get many lambdas on their premises, we will deliver them. However, a provider might well want to keep a provider relationship where the provider is the smart network provider and it can force you to work through it."

SURFnet policies are designed so that expenditures of the network serve as a magnet for other telecom players and communities to lay dark fiber. "Everybody in the Netherlands is allowed to dig fiber in the ground and own it. . . If one provider asks permission to lay a fiber, it will be announced and all others who are interested in the same route are free to use the same digging and then the providers have to share the costs of the digging together." The national policy of the Netherlands government is to create conditions such that the combined actions of the Research and Education and commercial sectors will create a fiber based infrastructure for the entire nation and will keep it in the top ten in European information technology infrastructure national expenditures.

State of the Internet: Light IP & Gig E, pp. 9- 10

From the COOK Report Annual Report: There appears to be a choice of two paths to our telecom future. One is to go with the highly innovative pure Internet approach of gigabit Ethernet over condominium fiber. Such a choice empowers the customer, facilitates decentralization over centralized control and provides small and innovative businesses with the environment that they need in order to flourish. The other path is to try to fore stall the innovation by squashing competitors with a massive vertically integrated company founded on older technology and leveraging access to content and over a network monopoly so pervasive that people will find they have no choice but to buy it. What could be in store for us all, if things go in this direction, was summarized by Scott Cleland CEO of the Precursor Group on Friday January 19th, 2001. "Precursor believes AOL-TW has budding 'Microsoft-like' potential to grow increasingly dominant by being the leading national company that brings together the various online interfaces (content, Internet access, buddy lists, instant messaging, etc.) to become the de facto consumer online access market standard much like Microsoft Windows brought together the various desktop applications to become the de facto consumer software market standard." See http://cookreport.com/lightipgige.shtml

Grids, pp. 11 - 18

In the context of broadband networks, grids are becoming a much discussed subject. We interview Peter MacKinnon who is Managing Director, Synergy Technology Management. Grids are seen as a pervasive computing fabrics into which users can plug. Says MacKinnon: 'computational grids' are viewed as a network of distributed computing resources that can work both cooperatively and independently of each other. They allow applications to operate in a distributed multi-platform environment across various geographic scales defined by the physical networks involved. Computational grids represent one of the frontiers of computing. They raise many fundamental challenges in computer science and communications engineering, much of which has to do with partitioning a problem across multiple machines, latency in the networks and administration and allocation of the grid's resources.

Then there are access grids. That's another way to look at the grid, where it provides access to devices, such as, say, radio telescopes or optical telescopes. I make these references in technical terminology or scientific terminology, because this is the locus of this grid frontier. It's not in the commercial world yet. For the time being we're not talking about applications that relate to customer relationships.

A more complex example called 'Neptune a Fiber Optic Cable to Inner Space' is actually a proposal right now. It's being led by Woods Hole and JPL in the U.S. and the Department of Fisheries and Oceans in Canada, at this stage. There are other players. Basically the intent is to put fiber optics in a grid-like form on the Juan de Fuca plate off the West Coast of North America. This will be an ocean floor-based grid that will have nodes spaced a hundred kilometers apart and a receptor. And in that receptor or junction box, which would be analogous to a satellite system or a space station system into which you'll be able to plug in instrument packages.

A number of financial systems could be turned into a grid by being connected within a single phone network. With these grids a global trading house, for example, could end up having its neural network system in London connected with that in New York, with that in Tokyo, with that

in Sydney. So now you're monitoring on a different level. If you have the latency problem solved, so that you can both do this in real time and do the computations required, it becomes a potential example that could lead to the development of new types of financial instruments or new ways of hedging or new risk reduction capabilities. The financial area would probably be one of the first major commercial uses of grid-like capability. Another example of grid-like capability could well be in utilities. These are organizations that have distributed systems already that want to use the grid, in a simple sense, to do status checks, self-healing, monitoring, or whatever the case may be.

The Grid Forum is basically a place to talk shop for everyone from those who are trying to build grids to those who actually want to use grids. It is a common meeting ground as a place to discuss the entire dimension of grids. They have organized themselves into several working groups. See http://www.gridforum.org/

Many of the technical issues needing to be solved involve integrating current communications and computing advances with the architectural needs of grids. However we have to find ways of solving certain problems in both areas before we are going to be able to make grids, in what we might call a promise, the success that they might appear to be.

For delivering the promise of grids, as demonstrated by the notion of grid space, which we just talked about, the following are likely. Grids will provide powerful, interactive, dynamic and flexible environments allowing for opportunities to create new discoveries, on a level of Grand Challenges. They will also allow for more R&D without increasing other resources as well as widen access and enhance educational uses.

Actual implementation of robust grids is going to require a great deal more advancement in software systems than is currently available. When you start to think about what it is that you're going to do in a grid-like problem-solving environment, then there are some really fundamental technical issues that need to be addressed.

If you have a computationally intensive problem, there have been a lot of advances made in development of parallel computing approaches in the last several years. These advances allow problems to be partitioned so that multiple parts of the problem can be simultaneously computed on different processors . Therefore, you have to understand in some detail the kind of problem that you're dealing with in order to know how to partition the processing.

Chicago Civic Net pp. 19 - 21 and RFI pp. 22-24

We interview Joel Mambretti Director of the international Center for Advanced internet research at Northwestern University about Chicagos plans for CivicNet, a public private partnership that would link all public schools and libraries and city agencies by fiber.

The City is doing what any large organization should naturally do, that is to say, it is planning for a future rise in demands, and it is planning for new services, while understanding that it has a set budget.

The projected CivicNet budget is $250 million over 10 years. The question is: How can we best optimize this expenditure? Part of this process is simply doing what organizations should normally do. Certainly, then, they want to ensure that their requirements are clearly communicated to potential providers of those services, because the providers are very eager to know what customers require so that they can respond. Another part of this process is establishing an ongoing dialogue between the City and potential providers of services to match what is being asked for with what can be provided. This is a process that is healthy for both sides and one that both sides appreciate

Therefore, the idea isn't to go out say: Build something. But rather to say to the general world: Here are the requirements, and then ask for a response. That is why the city has issued an RFI not an RFP.

Thanks to subscriber Frank Coluccio we received a pointer to the Civic Net web site that has been built by the city of Chicago.
http://www.chicagocivicnet.com/civicnet/SilverStream/Pages/civiclist.html The site and the RFI document itself – some 130 pages in Adobe PDF format - begins to bring home the seriousness of the project. For a project of this scope, it seems likely that the city is making unprecedented use of the Internet to bring a public focus to its civic net acquisition procedures. In addition to the RFI the web site contains a forum for respondents to question the city and it contains access to city mapping tools and map sets that in non electronic form would be hugely expensive to enlist in a venture such as this. Finally it is perhaps the most significant example of CANARIE's telecom outlook influencing events in the United States. We republish here a shortened form of the RFI text that is also lacking its extensive appendices. The complete document is worth examining in order to get a sense of how ambitious this project is.

Senate Looks at ICANN, pp. 25 -26

When the Senate Commerce Committee announced ICANN hearings on the heels of house hearing and invited At large Board member Karl Auerbach to speak, we were encouraged Unfortunately the Senators were not well informed. While Auerbach's and Froomkin's testimony was accurate and should have caused rapt attention, Mike Roberts and Roger Cochetti as two of the people most responsible for the mess sat there and made apologies for an ICANN that really was not yet mature and had not yet had a chance to do what it was put into place for. Senator Burns in the opinion of an observer called the hearings only because he figured if his House colleagues were concerned he better find out what this was all about.

The high light of the morning came when Senator Burns asked Michael Froomkin what he would propose in light of the criticisms of ICANN that were delivered to the Committee. Froomkin: "the most important issue is not setting a precedent by which a Department, like the Department of Commerce can end run the Administrative Procedures Act. And that is an issue that frankly is bigger than the Internet. The global concern here is not just in this process. For this is a way in which agencies can by pass ordinary procedures to create a privately organized regulator in all but name. A regulator that uses control over a federally dominated resource to make people sign contracts with it, pay it money and do what it says. And then not be subject to due process. Not be subject to court challenge. Not be subject to ordinary oversight. That is really cutting the congress and cutting the American people out of the regulatory process. So while in the case [of the seven selected new TLDs] you might have had an outcome which was better than no decision at all – I have nothing against any of the winners here. I have no reason to believe that any are bad or imperfect, and for all I know we would be all better off if they were all put in the root along with lots of others too. It seems to me that there is a good government issue that is pretty serious here. Someone needs to hold Commerce's feet to the fire on this one." Burns pledged a follow up hearing at to look at "redress of due process." Whether he really understood remains to be seen. As we have said before and as one of the other witnesses pointed out, if ICANN succeeds there will be other ICANNs all designed by corporate interests to engage itself dealing and ignore all due process rights of those whom they would regulate.

Testimony - Auerbach pp. 27-30; Froomkin, pp. 31 -35, DNRCI pp. 35 -37

Auerbach: There are those who say that ICANN is merely a technical body. I am a technologist. Yet I have a difficult time understanding how any of ICANN's decisions concerned with the Domain Name System have any technical content at all.

One must wonder where the technical component might be in ICANN's Uniform Dispute Resolution Policy - a policy that expands the protection of trademarks to an extent not granted by any national legislature. And one must also wonder where the technical component might be in ICANN's preservation, indeed in ICANN's extension, of the hegemony of Network Solutions over the naming systems of the Internet. We know more about how the College of Cardinals in Rome elects a pope than we do about how Icann makes its decisions.

There are lessons to be drawn from ICANN:
- ICANN has shown us that governmental powers

Continued from page 39

ought not to be delegated to private bodies unless there is an equal obligation for full public participation and public accountability.

- ICANN has shown us that a public-benefit and tax exempt corporation may be readily captured by those who think of the public less as something to be benefited than as a body of consumers from whom a profit may be made.

- The role of the US Department of Commerce in ICANN has shown us that Internet may be used as a camouflage under which administrative agencies may quietly expand their powers without statutory authorization from Congress or the Executive.

Froomkin: If in 1985 the Internet itself had been a proposal placed before a committee that behaved as ICANN did in 2000, the Internet would have been rejected as too risky. Risk aversion of this type is antithetical to entrepreneurship and competition.

Worst of all, ICANN applied its criteria arbitrarily, even making them up as it went along. The striking arbitrariness of the ICANN decision-making process is illustrated by the rejection of the ".union" proposal based on unfounded last-minute speculation by an ICANN board member that the international labor organizations proposing the gTLD were somehow undemocratic. (That this same Board member was at the time recused from the process only adds to the strangeness.) The procedures ICANN designed gave the applicants no opportunity to reply to unfounded accusations. ICANN then rejected ".iii" because someone on the Board was concerned that the name was difficult to pronounce, even though the ability to pronounce a proposed gTLD had never before been mentioned as a decision criterion.

DNRCI: The sad fact is that ICANN has been "captured" from the beginning. Special interest groups have dictated the direction of ICANN, and have morphed it into an Internet Governance body with none of the protections afforded by governments.

As currently constituted ICANN has failed on all charges. It has moved slowly; been unrepresentative; acted to limit competition; and failed to offer useful, fair, coherent policies, or even policies which encourage investment in virtual property. ICANN is a policy experiment that has failed.

ICANN is correct in that its formation was an unprecedented experiment in private sector consensus decision-making. Unfortunately, that experiment is in the process of failure. ICANN's claim of "openness and transparency, based on Internet community consensus, bottom-up in its orientation and globally representative" is far from the reality of the situation. ICANN is the classic top-down organizational structure without accountability. When its by-laws are inconvenient, they are changed without discussion.

The Internet is the single most significant communications medium ever created. Its power goes well beyond that of shopping malls and e-commerce, and empowers individuals in a way never before imagined. It is thus a national as well as an international resource. The ability to control important aspects of this technology cannot be underestimated. It is up to all of us to remain vigilant when organizations are given special privilege by a branch of the US Government to control this vast means of expression. Safeguards must be put into place whereby individuals, non-profit entities, churches, tribal governments, and other disenfranchised groups may provide unencumbered input and opinion to an open, transparent and accountable entity. This entity is, unfortunately, not ICANN in its current form.

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

Gordon Cook, President
COOK Network Consultants
431 Greenway Ave
Ewing, NJ 08618, USA
Telephone & fax (609) 882-2572
Internet: cook@cookreport.com