




# The COOK Report on Internet



© Cook Network Consultants

## Optical Border Gateway Protocol is CA\*Net3 Experiment To Enable Peering Without Purchase of Transit Through Carrier Cloud Direct Connection of Lambdas Could Reduce Backbone Transit Traffic Enabling Internet to Cope with Gig E Bandwidth Explosion

**Editor's Note:** Bill St Arnaud is the director of network projects at Canarie. Since 1998 he has been responsible for the coordination and deployment of the world's first national optical research and development network, CA\*Net3.

**COOK Report:** It has been nearly three years since we last talked. Most of the predictions you made ten have come true and the progress made has been stunning. Gigabit Ethernet and 10 gigabit Ethernet are going to bring an explosion of bandwidth at the edges of the Internet. But if traffic at the periphery of the network is going to grow by three and four orders of magnitude, people are wondering how backbone traffic can possibly grow by four and five orders of magnitude.

But first, what is thought to be the limit for speed increases at the periphery of the network? Are we ever likely to see one hundred gigabit Ethernet?

**St Arnaud:** One hundred gigabit Ethernet in serial mode is rather unlikely. When you get to ten-gigabit speed in serial mode, you begin to experience all kinds of non linear effects in the fiber. For example you get polarization dispersion and Raman scattering. Consequently, while I do not see one channel at 100 gigabits on one strand of fiber, I do see the possibility of 1000 channels of ten-gigabits on that one strand of fiber. One company recently announced that they have developed a Bragg grating that will support up to 880 channels on a single strand of fiber. Although this may have been a "hero" experiment in the lab, it does point to the future direction of fiber capacity.

There are a number of ways to multiplex optical channels on a single strand of fiber. One is called a Bragg channel. Another is called an optical wave guide. There are a

number of others. They are competing with each other in a search for the best ways to multiplex all these optical channels together.

**COOK Report:** How then would you pose the current problem? We are having an explosion of bandwidth at the edges of the network. We are having enterprise networks building their own private networks out of gigabit Ethernet. How then would this one group at the edge of the network communicate with another group somewhere else on the edge of the network with a group that could only be found by transiting a backbone to what we might call the "legacy" SONET-based Internet.

**St Arnaud:** To answer your question we must examine the current context of data networking. We are witnessing a profound upheaval in the traditional customer carrier relationship. The carriers look at building their networks from the same old standard of a centrally controlled network. Consequently, they tend to be able to see the only solution to the problem of their customers' growing traffic as being one of just adding more of the same kinds of SONET channels that are already there. They then turn around and offer to their customers these new channels as managed SONET services. The major challenges to this approach of networking is scalability.

**COOK Report:** When you say carriers, would you make any distinction between the legacy carriers like Bell Canada, AT&T, and Sprint and the new green field players like Level 3, Global Crossing and Williams?

**St Arnaud:** Rather than divide the field into old and new carriers, I would prefer to talk about two different models of building networks. The first I would call the centrally controlled command economy and the second what I would call the capitalist or more

Volume IX, No. 8, November 2000  
ISSN 1071 - 6327

free-market approach to networking. Most networks today are designed and operated through a "command economy" architecture.

### Two Different Ways to Build Networks

You have a single carrier with a large and sometimes even global vertically-integrated infrastructure where the carrier is accustomed to providing services at all levels of the protocol stack. Such a carrier does everything. It deploys the fiber. It lights the fiber. It adds SONET services to the fiber. It provides ATM services and on top of that adds E-Commerce and web hosting services. Such an integrated carrier finds itself running a large and very complicated bureaucracy and also finds that it must have a large and expensive management system to provide and integrate all of these complex and varied services and activities. Some people are beginning to ask whether these huge integrated a vertical carriers can continue to scale.

Let's call the other approach to building net-

#### On the Inside:

Optical BGP	pp 1 - 9
World Wide Packets	pp.10 -14,26
IPv6 Inevitable?	pp. 15 - 18
ICANN Election	pp.19 - 24
Executive Summary	pp. 25

works the “capitalist economy” approach. With this approach different companies perform the different tasks necessary to making the over all network function. This is a much more horizontal form of networking and one that allows much more control to be exerted from the edge. While some users may say that they want to have “cradle to grave” network services other users may say I want to be able to route and control my own wavelengths independent of a central controlling authority. The company that provides the fiber, may be separate from the company that provides the wavelengths which may be separate from the companies that provide Internet and other advanced services. Moreover I want to be able to connect to my preferred Internet service provider who may be halfway around the world.

It is the later form of “capitalist economy” network architecture that will be enabled by all these optical channels. When you can get a thousand wavelengths on a single fiber you can sell individual wavelength to individual customers. Your customers may even be able to bypass the traditional transit supplier by constructing their own wavelength network.

**COOK Report:** Are you saying then that if you are one of the large installers of dark fiber like level 3 or Williams or Qwest, you will have a variety of solutions at different levels of complexity and cost for your customers. Whereas if you are working with an older, more mature service provider like AT&T, MCI and Sprint, you will find that you will have an obstacle in migrating your own business in this direction?

**St Arnaud:** I think that it will always remain a mixed market. I think that you will have customers out there who will pay a premium to have “cradle to grave” managed Network Services. These customers do not want to manage their own wavelengths and they don’t want to learn how to run Gig E over dark fiber. They are willing to pay a premium to all this taken care of for them. These folk still form a large part of the market and it will take a long time for these things to evolve. The market then will evolve along with the carriers, some of which will die and some of which will become successful.

Right now there are thousands of carriers worldwide and yet we only have perhaps 50 different automobile companies worldwide. I think we will see a very significant aggregation of traditional backbone carriers in this new marketplace. Those who adopt quickly in order to take advantage of new technologies and economies of scale will survive.

**COOK Report:** In evaluating the carrier market what do you look for to distinguish those who will survive?

**St Arnaud:** Those carriers who partner with the research and education community I think will be the survivors. I believe that you will begin to see considerable innovation in developing new network technologies and business models coming from the research and education community. Five years ago most people believed that the research and education community’s role in the development of the Internet had diminished to the point of insignificance. The assumption was that the private sector would be successful in pushing the rollout of broadband Internet services to businesses, schools and homes.

**COOK Report:** Why didn’t this happen? Are you in the research and education community in possession of a flexibility that the private sector business market does not have?

**St Arnaud:** Perhaps because the university research and education community with high bandwidth demands can afford to be an early adopter of new technologies that production networks cannot afford to experiment with. These institutions are now partnering with the private sector so that both communities win in the process of building next generation dark fiber networks.

[Editor’s Note: On the NANOG list on September 1 Sean Doran answered a question about Internet 2 (the U.S. next generation network) by saying: “I’m sorry, but parts of my network [Global Telesystems] are ALREADY faster than Internet2’s network by a factor of four. Other fully commercial networks are also well beyond the technology that most academics brag about. The ONLY interesting large-scale academic networking project is currently being done by CANARIE. “We agree with Sean’s assessment. Bill St. Arnaud is the Director of the CANARIE Ca Net3 effort.]

## Public Private Partnerships

**St Arnaud:** What we have seen in Canada is that this partnership led by the university research community benefits everyone. In Quebec, for example, the universities are building a fiber network some 3,500 kilometers long that will connect all of their institutions. The cost for installing the 3500 kilometers of fiber comes to less than 2 million US dollars. This is done as a partnership with five or six new carriers. By contributing to the cost of laying the fiber, the outcome is that, at the end of the project, each of these new carriers has access to a network as powerful as that of the incumbent telco.

**COOK Report:** Who are these new carriers?

**St Arnaud:** They are Cogeco., Gateway

Telecom, GTE Telcom, Dixon Fiber, QuebecTel, Videotron and VDN. You wrote about VDN and its activities in Montreal in your last issue.

**COOK Report:** I have heard the thought expressed that in some of the research and education community in North America high-speed networking is seen more as a means of soliciting subsidies to pay down the campus cost of connection than as a way to develop these more innovative industrial partnerships. Are you willing to comment on this opinion?

**St Arnaud:** Yes I think there is a mixture of interest out there with some of the University CIOs playing a vital role to help their respective communities to get advanced Internet services. Others are interested just in their own internal needs. Perhaps a university in a big city where there are many other universities doesn’t feel the need of being a strong community player. But others like the University of Indiana or the University of Illinois have been playing a very strong role model in their respective communities.

**COOK Report:** Are you aware of the emergence of many new carriers in the U.S.? Of course we all know about Qwest, Level 3, Williams, Global Crossing and similar large players. But who else?

**St Arnaud:** You certainly have many new CLECs and DLECs. You have many more of these in the U.S. than we have in Canada. I would keep my eye on a new carrier called PF.net. They have a new management team that I believe who truly “get it”.

**COOK Report:** And by DLECs you mean?

**St Arnaud:** A CLEC is a competitive local exchange carrier focused on the voice market where a DLEC is the same thing only focused on the data market.

**COOK Report:** What then is the implication of all this growth in traffic and data at the periphery of the network for what one might call the first generation of the Internet backbones that grew up and matured between 1995 and 1997? Given what is happening now, you have the old hierarchical question of who gets to peer with whom, and a relatively small number of backbones facing a collision with the huge amount of data being generated at the edges of the network. How do we cope?

## Internet Backbones: Changing Demands

**St Arnaud:** You’re asking precisely the right questions. I think that the first generation of Internet backbones were represented by

Sprint, MCI, UUNET and GTE Internetworking. One might also refer to these as the first business model for Internet backbones. The second generation or Business Model is represented by Cogent Communications whom you interviewed in your September newsletter and I believe the third generation or business model is a back plane built from customer controlled lambda's or wavelengths of light.

For example, right now if I were a small ISP in upper New York state, I would usually peer with an upstream provider or two that had POPs close to your network. For a fee, these large upstream providers would transit my traffic to the rest of the Internet. But if I am an ISP and I can control my own wave lengths, I have the possibility of bypassing or routing around a transit provider.

If ten percent of my traffic were going to and ISP in Hollywood California, it might be cheaper for me to buy a wavelength of light and route around a traditional transit supplier by taking my wavelength directly to that ISP's border routers on a no cost peering basis. I think that the availability of these wavelengths where you can peer directly with other smaller ISPs will allow you to build what I call octopus networks. In such a network you take the central part of your network and then you extend all these tentacles to other networks with whom you share significant amounts of traffic. If necessary such a tentacle could extend across the entire continent and enable you to peer directly with a like-minded network. As wavelengths become cheaper and more prevalent, at some point in the future it will be cheaper to buy a wavelength rather than the equivalent transit bandwidth.

**COOK Report:** That is rather fascinating. A year or two ago if I were a regional network in New England or the Northwest of the United States or somewhere else, the economics of the business model would dictate that I should look for an exchange point into which I could bring my network. I would then peer there with as many similarly geographically situated networks of my size as would agree. My objective would be to do as much no cost peering as possible and drive down to the absolute minimum the amount of traffic going into and out of my network that I would have to pay a transit provider to carry.

What I think I hear you saying now is that there is technology developing that would enable smaller regional ISP is to achieve these goals without taking leased lines from the borders of their network to several network exchange points. Is this correct? And if this could we explore how this technology will work?

**St Arnaud:** Yes. The current business model

is for ISP to go to 3 or 4 local exchange points and then buy transit for the bulk of its traffic from one or more upstream ISPs. But if it can purchase affordable wavelengths, it can connect to many more exchange points or by privately interconnecting and directly peering with many more ISPs.

**COOK Report:** Then this does not cut exchange points out of the picture. It gives ISPs many more options and very likely less expensive options for interconnection with the rest of the Internet?

**St Arnaud:** That is correct. Currently most commercial DWDM fibers have anywhere from 32 to 40 and up to 80 wavelengths on them. DWDM systems supporting over 800 wavelengths are being tested and it seems very likely that within a year or two, a single fiber will be able to support wavelengths numbering to well over a thousand. With such a proliferation, a wavelength should be cheap enough so that an ISP in New York could easily afford to take a high bandwidth connection to the PAIX exchange in California rather than purchasing a transit service from a major upstream provider.

**COOK Report:** What we are talking about then presumably is optical BGP which you have proposed as an extension to the border gate way protocol. If so, please explain in as much detail as possible how these extensions will work.

## Optical Border Gateway Protocol

**St Arnaud:** the Optical Border Gateway Protocol (OBGP) is an experimental concept which at this point is unproven. It is very radical and there may be many potholes that we encounter on the way to implementation. The first enabling step is dark fiber and the availability of many dozens of wavelengths such that an ISP can then purchase its own wavelength and use it to connect to an exchange point several thousand miles away if need be. At such an exchange point you will be able to do standard peering. The advantage given you by OBGP is that you can get to that exchange point without having to pay a transit costs to an upstream provider.

As the number of available wavelengths multiplies and the prices those wavelengths comes down, ISPs will find that they are able to buy wavelengths from their networks to dozens and perhaps even hundreds of exchange points some of which will be in other countries and on other continents halfway around the world. This is the first step. It's doable today and is beginning to happen as Williams and others are now selling wavelengths to ISPs.

The next step is where OBGP comes in.

Rather than buying a wavelength to a fixed place like the PAIX in California one could buy a wavelength to an appropriate optical switch. Each optical switch in the path to the PAX, in fact, then becomes an exchange point. With OBGP any optical switch in a optical network can be treated as an internet exchange point such that autonomous ISPs can interconnect and peer with each other anywhere along the network - in essence a truly distributed IX.

Now imagine an ISP in California that wants to peer with an ISP in Florida. Further imagine that both of those ISPs own wavelengths that travel to or through the same switch in Salt Lake City. With OBGP that ISP in California, which owns a wavelength to the switch in Salt Lake City, can see your wavelength coming in from Florida to that same switch. Consequently, you will be able to route traffic to and from the California ISP's wavelength without any intervening AS number or network control by a central authority. You will, in effect, have established no cost peering with that ISP at the switch in Salt Lake City. Furthermore, if another ISP in Colorado connects to that same switch, your can reroute your wavelengths and re-establish your peering session to the Colorado ISP instead of the Florida ISP. The beauty of this is that you don't need a router in Salt Lake City - only control of the optical ports. But for all intents and purposes, to your peers, it looks like you have full fledged layer 3 IX in Salt Lake City.

It is important to note that many of these concepts originated with ATM. What we are doing is stealing some of the best ideas out of the ATM world and converting them to the optical layer. But instead of a traditional "central managed" approach to networking as is currently being planned for optical networks, we are allowing the users at the edge control the routing and management of the wavelengths.

Now this has some profound consequences. If you let users at the edge control the routing and topology of the network that means the carrier in the middle will be very limited in how it can optimize and manage the wavelength routing in the network.

The analogy I like to use is building railways versus highways. Railways are operated very much like telecommunication networks today - the same organization that owns the infrastructure also controls the routing and scheduling of the trains. However, with highways the organization that builds and maintains the roadbed has little control over the routing and scheduling of the automobiles and trucks that use the highway. The user makes the choice on what interchange they want to get off and on the highway.

**COOK Report:** This is very powerful stuff.

Who is behind OBG? Where is it being developed? What needs to be done in order to finish it?

## The OBG Testbed

**St Arnaud:** First of all I must emphasize that this is a very experimental concept. At this stage it is high risk and there may be some kind of fundamental flaws that we haven't thought of yet. The background event for the OBG effort was the dark fiber build in Ottawa to the various universities and research institutes. All these organizations speak BGP in their networking environment. Moreover they are usually multi-homed. The only problem is that it is not practical to build a total mesh of fiber to every institutional user. We have 26 different sites and laying down a full mesh of fiber between the sites would require thousands of different fibers — something that clearly is impractical.

For example we have dark fiber going from our offices here in downtown Ottawa out to a research facility on the west end of the city where a number of fibers terminate. Also a number of fibers terminate in the east end access point. Now I would like to be able to route my traffic to any one of the institutions to which I do not have a direct fiber connection. One model of doing this is to put a carrier in the middle. The carrier would say "OK I'll take care of it. I can build this cloud whether it be ATM or one that is made out of wavelengths. I will route your traffic and don't you worry about it."

When it comes to optical networking the carriers can't conceive of anyone besides themselves owning the cloud and doing the routing for their customers. But we found that owning the dark fiber ourselves seriously changes our perspective on the possibilities of networking. We no longer want to go back to the model where somebody has to be in control in the middle in order to aggregate and redistribute traffic. Customer owned dark fiber has given us freedom. We like the freedom and independence of controlling our own routing and never want to go back to having a carrier in the middle.

We then realized that one way to avoid this might be to have an optical switch in the middle and for each of us have access and control of our own separate ports on that switch. If you give each of institution control of ports on that switch, we can then decide for ourselves how to route among us and do so without the interference of any central organization. Now to look at how this might work in other situations, the dark fiber supplier may be the one who installs the switch to be used by his customers. In such a situation the fiber supplier provides the electrical power for the switch, leaving the control of the switches ports to his custom-

ers. He says to his customers: "you cross connect to whomever you want in that switch and we don't even want to have anything to do with managing the switch expect providing electrical power"

**COOK Report:** So the provision of these services would be an attractive business model for a company like Williams?

**St Arnaud:** Yes as well as for Metro media Fiber and for companies like Equinix. This is a service that would be natural for anyone running an exchange point to provide, particularly for those who don't want to function as a full services provider. They just want the minimum services of the electrical power and a secure place to house the switch on which you depend.

**COOK Report:** You are looking at OBG as a means by which to enable these kinds of interconnections? When you say you are uncertain of the feasibility at this point, what do you mean? What kind of obstacles might you run into? Might they be technology or the expense of integration with existing software and hardware?

**St Arnaud:** As of today we don't see any such obstacles. But for every great idea that succeeds, there are a hundred that fail. It could be market reasons, or business reasons, or there could be a technology component that we've so far overlooked. When we build a prototype as proof of concept, we will have a better feel of whether or not it is really workable.

**COOK Report:** Are you working on a prototype in the Ottawa area or somewhere else?

**St Arnaud:** Exactly. I have the switch sitting right here in front of me. We hope to have the first prototype in operation later this fall.

**COOK Report:** What do you have to do between now and then to get it working? You said that that is merely an extension to the border gate way protocol. Is the only thing you have to do to write some new code and add that code to working BGP software?

**St Arnaud:** Yes. That is why we are and 99.99 percent certain that it will work.

**COOK Report:** What then is an extension to a protocol? Exactly how you create such a thing?

**St Arnaud:** In most protocols you have an options field. BGP has such a field and a number of proprietary products have been using that field for their own special purposes. We are proposing to use the options field in BGP to turn networking upside down. Today when you connect an ISP to an upstream ISP, the first thing they do is to

install the physical fiber. Then they put in the link layer which can be ATM or SONET. Next they establish IP connectivity and finally BGP connectivity. We are saying let us reverse that whole process.

So first we would establish BGP peering. I would say I want to peer with Gordon Cook. I have your AS number and to start a BGP session where I instruct my router to connect us. It is the router then that establishes the physical connection between us. Right here in Ottawa most of our institutions are going to put in either for four or eight wavelengths. They can have their wavelengths set up to go to whichever universities they choose. Then if someone says let's set up a BGP session with Gordon Cook, a router will take one of those wavelengths and steer it towards a connection that will peer with your network.

**COOK Report:** Does bringing OBG to market depend on developing a common means of modifying the options field in BGP? And therefore you just need to develop a common way of going about this so that everyone operates with the same expectation.

**St Arnaud:** Yes. The first thing we will do is a proof of concept that will demonstrate these capabilities in the next couple of months. We will put together and test a very simple version of the protocol. We already have a number of industry partners involved in the project who have indicated to us that they're ready to take the concept commercially and go with us into the IETF.

**COOK Report:** You will have "bof" and then form a working group to make this an official IETF standard?

**St Arnaud:** That's correct. But we don't want to actually start on standardization efforts until we are certain that it is workable and that there are no unexpected "gottchas" in the project.

**COOK Report:** How complex is the final product that we're looking at? Is it a document of 10 pages or hundred and 50?

**St Arnaud:** To anyone who is knowledgeable in BGP it is a relatively minor change and could probably be done in only a couple of pages. It is not a big radical new protocol but rather a very simple modification to an existing one. The bigger challenge is in modifications to the routing software itself. In our proof of concept project we are using public domain software. Once we have successfully completed the proof of concept project, we certainly hope that Cisco and Juniper and others will modify their own router operating system software to incorporate our optical extension to BGP.

**COOK Report:** So much of what you need to do is to get everyone to agree to an appropriate way to modify the options field in BGP. Once you have done that Cisco and others will have to change their running software so that it will recognize optical extension data from this particular options field?

**St Arnaud:** Yes. And then the commercial vendors will also add their own set of bells and whistles — management software and things like that. They will have to see value in what we've done and have to be confident that there is a market for adding our extension into their software.

**COOK Report:** Why wouldn't they think there is a market for it?

## OBGP Runs Against the Accepted Grain

**St Arnaud:** Right now in the Optical Internet Forum and in the ODSI working group the traditional industry view is completely opposed to this concept. The traditional equipment suppliers are working on a completely different strategy. In the traditional centralized, carrier-centric view of the world, the paradigm is to take the ATM model of networking and replace the ATM virtual circuits with wavelengths or lambda's. In this model the carrier totally controls everything.

**COOK Report:** this is what Francois Menard meant when he said that Sycamore was going to do circuit switching in colors?

**St Arnaud:** That indeed is what they're talking about. There are two approaches to traditional optical networking. One is an overlay model where you have a circuit switched optical network and run IP on top of it. You would do this like you run IP over ATM today. The other approach is called a peer networking approach where the wavelengths are treated as MPLS Logical Switched Paths.

**COOK Report:** What is the advantage of doing it this way? Is it being pushed by the carriers who have big bucks?

**St Arnaud:** Yes it is something that the carriers who have big bucks are very comfortable with. The big carrier market is a huge market. If you want to sell to them and make a lot of money from doing so, you better deliver the products they want. Therefore the reason that the optical companies are focused on this approach is that they expect to be selling to carriers for the next hundred years just as they've sold to them for the last hundred.

**COOK Report:** So what then are you doing with all of this in Chicago?

**St Arnaud:** We are working closely with Joel Mambretti, Tom Defanti and Charlie

Catlett and other people at Internet 2 on these concepts and on the development of the OBGP protocol. They are working on building a dark fiber network in Chicago. The network will hopefully include the first optical network exchange where we will put these types of protocols and technologies to use. They are working with our Canadian industry and research partners on building the optical border gate way protocol.

**COOK Report:** If you look at carrier's attitude where they want to do circuit switching in colors, do you think it safe to assume that this will happen?

**St Arnaud:** Some form of circuit switching will continue on its own for a considerable period of time with perhaps much smaller proportion of the telecommunications market doing it this way.

**COOK Report:** At the same time, is it your belief or hope that you can begin to grow and develop the other approach dependent on OBGP? Do you see some possibility of the two approaches coexisting with each other for some period of time?

**St. Arnaud:** Yes and I think the circuit-switched approach will likely be overwhelmingly predominant for the next five to 10 years. This is due to the huge legacy infrastructure and the huge critical mass of equipment belonging to the existing carriers. We are at about the same development stage with OBGP that the Internet was in 1975. Because of the huge existing investment in the traditional models of networking, I think that OBGP will start with our university friends across Canada and with our colleagues in Chicago. It may evolve a little bit faster. I may be too pessimistic. But I don't think it will be mainstream for a number of years.

**COOK Report:** What will be the critical determining factors of progress? Economics? Inertia? Politics? What will be the critical determinant of what you see as a rather slow rollout and evolution?

**St Arnaud:** I think the speed of the evolution of deployment will be directly related to the speed of the rollout of dark fiber. When universities and schools and other institutions can get their hands on dark fiber, this becomes a very appealing course to take. Therefore the speed of the rollout of dark fiber will really be the big factor in this and not so much the protocol itself. In jurisdictions where fiber is readily available like in Chicago and Canada, you will see it happen relatively quickly. In other places it may take a longer time.

## Cost Implications

**COOK Report:** When I first started talking to you two and a half years ago you were

saying way back then that you saw two orders of magnitude in cost changes coming down the pike. in other words in the networking power that you could buy for \$100 in a new SONET based network you could buy today for a dollar in with the new IP over fiber networks.

**St Arnaud:** Oh yes. No doubt about it. Just look at the Cogent business model as an example.

**COOK Report:** Indeed on the level that you first enunciated it has been amply borne out. Now if you look at the new OBGP technology, do you think there's any likelihood of anywhere as near a significant cost savings as was gained with the new network technology after 1995. Could another order of magnitude of cost savings be added to the first two?

**St Arnaud:** Yes and no. First of all if you use dark fiber you get a dramatic savings in cost. For one time investment as little \$10,000 you can get your institution hooked up with dark fiber that will last for 20 years. It has gotten very very cheap. The biggest cost now for many of our institutions who have dark fiber is the actual Internet access itself. Our biggest single cost is the purchase of Internet transit service. Now if OBGP works out as we hope it means we can significantly reduce our transit costs. We will be able to peer directly with the networks of our choice and be able to reduce our transit traffic to say no more than 20% of the total traffic rather than the 80 to 90% that it is today.

**COOK Report:** I understand and this leads me to recall that while, on the one hand, we have a handful of extremely large ISPs like AOL and Earthlink, on the other hand we still have probably four to five thousand small ISPs in North America. It seems to me that these small ISPs better hope for the success of OBGP which should enable them to spend far less on transit and it remained more cost competitive with those who are far larger than they.

**St Arnaud:** Indeed. I think that is a correct assumption. What I think will be the first thing to happen first is that a large university with good networking skills sets, instead of spending all its money on transit services will be able to peer directly with many ISPs on a no cost basis. It will then expand from the university research community to large commercial enterprises.

**COOK Report:** I have heard comments about GigaPops. Is this where someone at the University goes out and builds a regional exchange point for that University and for local ISPs? Or perhaps will it be Equinix or some other commercial exchange point and the universities will come to that? And

if I'm an ISP in the local area, it would be to my advantage to connect and there as well. Right?

**St Arnaud:** Yes. You would interconnect your wavelengths there and try to have as many direct peering sessions there as is possible in order to reduce your transit costs. GigaPOPs and commercial exchange points seem to be the most logical places.

## Strategies to Cope With the Traffic

**COOK Report:** OK. If you look at how backbones are being run now, you have a situation where all the traffic generated by the explosion of gigabit and 10 gigabit Ethernet at the edges of the Internet, can perhaps not even be sent through all the backbones that we have now. How are you going to cope with it? How are you going to get it from one side of the network to the other? From a software and hardware technology point of view is it feasible to think in terms of three orders of magnitude or even an increase of four orders of magnitude in backbone speeds needed to balance the delivery of the even higher bandwidth generated at the edge of the net?

**St Arnaud:** Definitely one of the challenges facing us is whether the core routers on the backbone will be able to keep up with all the increase in bandwidth at the edge. Scalability is one of our biggest challenges. One way to get around the bottleneck of the router is to do what is known as cut through wavelength routing. If I have enough traffic to one destination my upstream provider may take that traffic and put it on a wavelength of its own and send it directly to its appointed destination rather than try to packet process it in its local core router.

**COOK Report:** This sounds like you are setting up a permanent virtual circuit but is what you're doing something different from what you see if you were using MPLS.

**St Arnaud:** yes. It's a matter of how is carried out. This cut through concept in both ATM and MPLS is well-established and well-known. The cut through routing is done by the carrier inside its "cloud." The customer is unaware of what's going on. As far as you know you are sending all your packets to a router and it's the router that makes the decision whether to cut through or not. But with OBGp you the customer gets to make the cut through decisions - not the carrier.

**COOK Report:** Does the existence of this OBGp technology create an opportunity for companies like Equinix who are building Internet exchange points to step forward and say to ISPs we have the kinds of optical

switches at our exchange points that will enable you to implement OBGp and we can assist you in learning how to accomplish these implementations?

**St Arnaud:** Yes we believe this is a very significant opportunity for organizations like Equinix or any of the other exchange point or GigaPOP operators. They would provide that neutral site and electric power for the switch and let their customers do their own cross connects using OBGp.

**COOK Report:** If you are looking at a backbone like UUNET are you really going to have to assume now that such a backbone is already a mesh of OC 48 and before long OC-192 permanent virtual circuits?

**St Arnaud:** I think that is a fair assumption. For example they may own a fiber from New York to Chicago on which they will run ten OC-192s terminating on a router at each end. Now they could have another 10 wavelengths on that same fiber and of those ten five might go through Chicago terminating in Salt Lake City and the other five go through Chicago and terminate in San Francisco.

**COOK Report:** With the carriers then, you will find almost exclusively SONET-based wavelengths because of their legacy of network equipment?

**St Arnaud:** Yes. Almost exclusively.

**COOK Report:** If I am a Cogent Communications and all the pair of fibers I can do my backbone somewhat differently from UUNET, can I not?

**St Arnaud:** The Cogent backbone is probably very similar to UUNET's. While they both use SONET framing they do not use SONET rings or SONET protection features.

**COOK Report:** But ten-gigabit Ethernet will include a wide area network framing as part of the standard. In about year or so when products using these standards are available, it should be possible to use 10 gigabit Ethernet framing in wide area network backbones?

**St Arnaud:** With the ten-gigabit standard you have three ways that you can do data transmission. You can transmit to it native as in a LAN with standard Ethernet frames. You can transmit data via a wavelength. For that option there is a slightly different form of the ten gigabit protocol. Or you can map your data to a SONET frame. Nortel, Lucent, and Alcatel, the traditional telephony oriented companies, are very strongly proposing that ten gigabit Ethernet be mapped to ten-gigabit SONET frames. Such a mapping will obviously leverage their massive market and investment in SONET.

**COOK Report:** How do you do the SONET framing?

**St Arnaud:** You take your 10 gigabit Ethernet frame and strip out the header and all extraneous information and repackage what is left into the SONET frame.

**COOK Report:** You do need a separate piece of equipment to do this don't you? And this piece of equipment adds cost and complexity.

**St Arnaud:** Absolutely. There is a standard called Ethernet over SONET (EOS). The big companies behind this are saying: let's map many Ethernet channels onto one SONET channel. If we do this, we will benefit from aggregation because people will not be using their Ethernet channels at full capacity. By multiplexing them together we can gain efficiencies into the use of a single SONET channel. But then you have to do buffering and queuing and essentially make yourself a router. Some people say why go with Ethernet over SONET? Why not use packet over SONET? At the present time there is quite a debate going on over these choices.

**COOK Report:** What would be a reason that you do packet over SONET rather than Ethernet over SONET?

**St Arnaud:** Let's assume that you're carrying IP packets and that you are carrying Ethernet frames over SONET. These Ethernet frames are adding only complexity and inefficiency. IP, in other words packet over SONET, is much simpler.

**COOK Report:** Can you deliver then IP over SONET to the edge of the network and right into the office building?

**St Arnaud:** Yes. But generally it is not done that way. Usually it is converted into an Ethernet frame at the local POP and transmitted over dark fiber to the customer. As a result some people say you might as well carry that Ethernet frame all the way across the network.

## Politics of OBGp

**COOK Report:** Let's go back to the question of OBGp. Assuming that this will work from the technology side of things, what kind of political economic issues could knock it off the track?

**St Arnaud:** Of course we believe that this is going to be the next biggest thing since TCP and that it will revolutionize the industry. But, in the past, many people have come up with similar visions that have not delivered on their promise. Therefore I don't want to be overly optimistic. We think it will have a major impact on the future of the network,

but the reality is that many other people have had similarly good ideas in the past that have ended up in the trash heap.

**COOK Report:** How much of the outcome either in the U.S. and Canada depends on the regulatory point of view that you take? Doesn't the predominant regulatory point of view focus on a vertical approach where the regulator looks at the relationship of huge vertically integrated companies to each other and permits the charging of one price for voice and another price for data packets. The question is does it make any sense to do this in view of the Internet? Or is continuing to do this merely a carrier over from the regulatory world view of the legacy voice networks?

**St Arnaud:** This likely derives from the situation where a large percentage of voice is still carried over switch voice networks which are very much more expensive to run than packet networks are. When a large percentage of all voice is moved to the packet network, then you are absolutely right. At that point there will be no difference in cost between delivering the two services.

Say you are an ISP and that 20 percent of your voice traffic stays on your IP network. But you also have to connect to the PSTN for the remaining 80 percent of your traffic. The cost of this PSTN traffic will always be much higher to you, than the cost of voice traffic generated purely as Internet data. And voice traffic for the next 50 years probably will always be more expensive than data traffic.

**COOK Report:** Is it fair to assume that this will be influenced by a kind of regulatory inertia? You assume that there is going to be a difference in cost between voice and data. Because the cost is greater if you permit the companies to charge more for the old way of doing things, you're really subsidizing the continuation of that way of providing services. Are you not?

**St Arnaud:** Yes and no. The dramatic drop in long distance voice costs have not really been caused by the Internet but have been caused by competition. The carriers had to wring more savings out of their legacy networks or incur losses. Traditional voice networking is actually very very hard to do. Real time delivery of the human voice requires the addition of all sorts of quality control mechanisms to the network to avoid such things as echo and jitter. Building a good network for voice is very difficult.

Many companies announced that they will do voice-over IP and totally eliminate the telephone company. . Voice turns out to do very difficult to do especially when you have to interface with the traditional public switched telephone network with its SS 7

signaling. We simply cannot ignore the legacy PSTN, in building a voice-over IP networks. Until the last telephone and the last country in the world is switched over to IP, we will still have to interface with the old legacy system. This need to interface with the legacy PSTN means that we have to have complex of protocols and signaling and things like that which drives up the cost of delivering voice whether it is over IP or the PSTN.

**COOK Report:** If you look at voice over IP as run on large corporate virtual private networks, can you get quality in those environments equal to the quality you get with a very different and more expensive equipment on the public switched telephone network?

**St Arnaud:** Absolutely. This is happening today within many corporate private networks. Voice in the networks belonging to large corporations is being packetized and moved rapidly into IP Transport.

**COOK Report:** Say then that you have independent parallel networks — voice continuing on the legacy public switched telephone network and a voice-over IP on corporate networks. You cannot easily go from corporate network to the public network?

**St Arnaud:** That is the problem. Going between a voice over IP network and a circuit switched networks is very difficult. This is especially the case with the big carrier switched networks which have extremely stringent demands as a result of their SS 7 signaling. Carrier grade voice-over IP is a long way off and it is going to be quite some time before it becomes a cost-effective solution.

## Regulatory Issues

**COOK Report:** When American and European carriers buy new equipment for their networks, to what extent to do they buy the newer technology of the IP networks or are they almost forced into a continuing to buy the older technology that is compatible only with their legacy circuit switched networks?

**St Arnaud:** They buy some equipment designed to tie the to networks together but they've been rather slow to do so because of the extremely stringent reliability requirements imposed upon their voice network functionality by regulators. Because of the insistence on the traditional "five nines" of reliability, the traditional telephone companies can take up to five years to evaluate a new technology. Their time scales are on entirely different order of magnitude than those of the Internet where we are swapping out boxes every six months.

**COOK Report:** Is there a kind of the institutional dilemma and the setting a regulatory policy? Can you imagine a situation where if we had regulatory leadership with a "net headed" disposition and we could begin to promote the acceptability of a world where if you're IP over fiber service goes out you would have wireless as an alternative replacement for the time of the outage. Or will everything be governed by the inertia of the legacy networks where those people will make it impossible for such a radical change of priorities to take place?

**St Arnaud:** I'm afraid it will be the latter with incumbent IXC's and ILECs trying to milk all possible profit out of their legacy infrastructure.

**COOK Report:** Do you think there's any possibility that some say three to five years from now someone could begin to make the argument that because of the ubiquity of various telecommunications systems that you could rely on the ability to switch over momentarily from one system to another in case of an outage? Could this be justification for the promotion of regulatory policy in favor of doing just this rather than taking another 20 years to write off the investments in legacy SONET equipment?

**St Arnaud:** Increasingly they do. I'm sure the different approaches to these problems will be developed. However the other problem that traditional telephone companies have is traditional customer expectations. They expect their phones to work all the time. With 99.999% availability your phone is almost always going to be up and ready to reach anyone anywhere at anytime. In the Internet we have a different set of expectations and ways of evaluating reliability. Going from the culture in the telephone companies to the world of the Internet is going to take a long time. But focusing on a voice is focusing on the wrong problem. I say forget voice because voice will be a niche market within less than three years.

**COOK Report:** So one might think of the future as a group of parallel telecommunications services. Hear you would have legacy voice services going forward into the future, but you also have the Internet data services and voice-over Internet data services. you will have a migration of people from the legacy voice services toward the Internet and the only uncertainty is how rapid the migration will be.

**St Arnaud:** Yes, and I can see the two networks existing in parallel for 50 years.

**COOK Report:** Yes. But 10 years from now I can imagine that we might be paying \$20 a month for our always on, five nines reliable telephone service and \$80 a month for a suite of Internet data services including

voice-over IP, e-mail and Web access, and access to high-speed video and audio.

**St Arnaud:** Perhaps. Except I see these two figures reversed. The suite of Internet services may be \$20 a month and the legacy, but highly reliable, phone services \$80 a month because Internet data services will benefit greatly from economies of scale.

**COOK Report:** Well that is an intriguing thought because, if voice services are that expensive, it will add to the pressure to migrate away from them to the Internet data services.

**St Arnaud:** We think that there will be three parallel networks in existence. In fact we would suggest that policy not to try to converge all telecommunications into IP because doing so will increase the cost of IP. The residential telephone market will continue to provide delivery of telephone by twisted pair. There is nothing wrong with a voice-over copper which in fact works very well. A second network may well be broadcast video over coax which again works very well. There is no really compelling reason to deliver broadcast video over IP. I believe there will be a third network dedicated to only IP and this network can under some circumstances carry voice or video or both.

Your primary mechanism for receiving broadcast video will be either an antenna or coax. The primary mechanism for voice will be either a cell phone or your current copper phone. IP will include a voice-over IP and other value-added services but for the next 50 years I am don't see it as entirely displacing the other two networks.

## Fiber to the Home

**COOK Report:** What do you see as the economics of bringing fiber into the home? You have talked about bringing gigabit data services to every home in Canada by 2005. What will be the dynamics that drive what only a few years ago most people would have considered an unlikely outcome? Yet you now see it as possible and Worldwide Packets has built its entire business model on this as a premise. What will drive this?

**St Arnaud:** This is our greatest challenge and we believe this is where the research and education community in partnership with the government and carriers can really help to accelerate the deploying fiber to the home. For example. In Quebec a number of municipal governments and school boards have gotten together and put fiber into every public sector building for a one time cost as little as \$10,000 per building. This is a one time charge for the next 20 years. If these municipalities put in a dozen extra strands to the school, it becomes very easy for the private sector at minimal cost to run fiber to

individual homes from a node at the school. We are looking at the feasibility of partnerships to get public-sector buildings fiber connected in a way that will save them money. Once we do this, if we do it right, it is relatively easy for the private sector to extend that fiber to individual homes.

**COOK Report:** Do your politicians understand the importance of getting this technology installed as infrastructure?

**St Arnaud:** Yes, Increasingly they do so.

**COOK Report:** if you survey the scene in both the U.S. and Canada from the point of view of viable business models for getting fiber to the home, what do you see as viable approaches and who are the major players?

**Arnaud:** At the present this a small market, but it is growing rapidly. There are a number of companies in the area. Homefiber is a company doing it in the Palo Alto area and SBC is doing it in the Atlanta area. It is pretty well a slam-dunk to put fiber in new housing developments, especially if you have to build new roads and trench new sewer lines. The challenge is to get it into existing homes. In the Western US, WIN is a company that is doing it as well as several public utilities who are now building a fiber networks to existing homes and who have developed compelling business cases showing that this can be done.

**COOK Report:** Who are the experts in this area?

**St Arnaud:** World Wide Packets is one of the leading companies in this area. Another effort worth looking at is Grant County in Washington state where the state government leadership was critical. Washington State passed a law encouraging utilities to get into the fiber to the home marketplace. But they also wisely said we don't want you to be telephone companies. We just want you to install dark fiber and sell access to that.

**COOK Report:** if we look at the issue of scaling of Internet backbones needed to deal with gigabit traffic generated at the edge, perhaps we can see the existence of two paths? One path relies upon a mesh of permanent virtual circuits and MPLS and is sometimes referred to as the path preferred by UUNET and Mike O'Dell. Now would the other path to scaling be the one of taking light paths to exchange points and use of the optical border gate way protocol?

**St Arnaud:** It's hard to say. OBGp is too immature at this point make a difference. We like to think that in a few years it might be the dominant protocol. But that remains to be seen. Over the next couple of years I think that MPLS will be the dominant means of traffic engineering backbone networks.

**COOK Report:** is a third possibility the continued scaling of large point to point backbones?

**St Arnaud:** There are probably three possibilities. I am not a fan of the MPLS approach. MPLS was a technique designed to cope with shortage of bandwidth. Now with the ability to buy wavelengths of light on fiber bandwidth is not an issue. Consequently micro engineering the network is probably not necessary today. Also it does not look like we can build routers big enough to aggregate all the traffic. We will require optical "cut thru" or "bypass" circuits. The challenge that faces is who controls the cut through circuits - the carrier in the middle or the customer at the edge.

Now, doing it this way is not building a traditional circuit switched model. Circuit switched models imply that for every flow you must start a new circuit. The service will set up a switched circuit and send a web page to you, and then switch to another circuit and send a web page to someone else. This will not scale. Having direct peering with each other with wavelengths and doing bypass around other ones may achieve the same ends.

**COOK Report:** so the 50 to one hundred largest carriers of the world will be going to a large number of exchange points and these carriers might develop a model of relying upon cut through circuits to reach each of the 50 to one hundred other carriers from the exchange points. Yes?

**St Arnaud:** That is correct.

**COOK Report:** And then the third way, if it works, would be the optical border gate way protocol.

## Three Models of Networking

**St Arnaud:** That's right. So right now, in the standards bodies, you have three approaches. One is called the overlay and this is basically the circuit switched model. This is being promoted by under the label of ODSI. Another one is called peer networking. This is where all the wavelengths are treated like MPLS tunnels. The third approach is ours. In this we say let BGP be the controlling mechanism and let decisions be made more upon the lines of traditional Internet direct peering relationships.

**COOK Report:** Would you clarify the distinction that you are making between MPLS and circuit switched?

**St Arnaud:** Traditional circuit switching is tantamount to running IP over ATM. that's the ODSI approach. The wavelengths are

like ATM virtual circuits and they are switched all the time all over the place. In the ODSI model switching has merely been moved down to the optical level and the IP layer is unaware of topology changes in the network. The MPLS way of doing things is the cut-through model. In the ODSI world if something breaks, the optical plane takes care of it. In the MPLS world the MPLS router has more responsibility to do the cut throughs.

**COOK Report:** In other words does your routing take place at layer two or layer three?

**St Arnaud:** That's right. In OBGp is entirely a layer 3 or even a layer four model in which the BGP itself actually controls the set up of wavelengths.

**COOK Report:** And from a structural point if you, if what you're running is going to take place at layer two, does that equate to the only place where scaling is possible being the very large telco kind of model?

**St Arnaud:** OBGp is a model that we believe can scale the most because all the control is at the edge of the network. And why? Because each small group of people at the edge the network is able to come to their own decisions and implement or change these decisions very quickly. The other two models assume the existence of a huge central carrier in the middle. If at some point in the future you have millions and then perhaps even billions of wavelengths that have to be centrally controlled and monitored how this could possibly scale?

**COOK Report:** The see the MPLS model as the one that can transition to the control at the edges model of the optical border gate way protocol? Or do you see it as just another way of keeping central control?

**St Arnaud:** I think the intention is to rely on it indefinitely. After all there are very few people at the present time who know anything about OBGp. If you are in Nethead from the Internet World, the MPLS model of cut through is probably the one that is most prevalent and popular in the IETF. If you are a Bellhead from the telephony world the ODSI model of an overlay is more popular.

**COOK Report:** So in this sense your technology outlook is determined by the technology origin of its investors? Since it looks like these will be competing with each other for a while can we ask if there is any evidence for a distinction in the cost effectiveness of the two different approaches? Does the first one depend on SONET where is the cut through MPLS model doesn't?

**St Arnaud:** Yes and no to both. It is just that both assume a large organization exerting

central control. OBGp assumes that you have dark fiber and that there is no central control organization.

**COOK Report:** Is there any reason is that you can see over the period of the next few years where the cut through central control model might become dominant over the ODSI circuit switched model? Or do you think that for the next several years those two models will coexist and keep similar market shares?

**St Arnaud:** It's difficult to say because it is so new and still untested. But in my personal opinion MPLS will hold a far larger market share because it is more in line with the Internet view of the world, while ODSI is more in line with the traditional carrier view of the world.

**COOK Report:** To conclude — if you were to offer any advice someone who is making decisions about acquiring telecommunications infrastructure in their community or to someone who is making investment sessions, what would your advice be?

**St Arnaud:** Number one advice. Wherever you are. Get dark fiber. Doing so changes your whole outlook on life. It puts you in control and will start to open your mind to all sorts of new ideas and concepts. Imagine a world before electricity. If you never had electricity and I am trying to explain to you the concept of what a blender does, you simply can't understand. But when you have electricity your whole world starts expanding. And until you get dark fiber you really cannot adequately grasp the possibilities of OBGp.

---

## Alcatel Transmits Five Terabits on Single Fiber

**[From an Alcatel press release Sept 9]** Alcatel sets a new world record for DWDM backbone networks at 5.12Tbit/s and paves the way towards 10Tbit/s transmission on a single fiber Paris, 7 September 2000- Alcatel, the world leader in optical networking, today announced that it has set a new world record for DWDM (Dense Wavelength Division Multiplexing) transmission capacity over backbone networks. The company has achieved a unidirectional transmission of 128 channels each modulated at 40Gbit/s, reaching a total throughput of 5.12Tbit/s (5,120Gbit/s) over 300 km. This result was achieved with three technologies developed by Alcatel: 40Gbit/s DWDM systems, an innovating technique to make the most of existing bandwidth, and TeraLight™ optical fibers. This demonstration once again highlights the technological leadership of Alcatel in 40Gbit/s DWDM systems as well as the unequalled potential of the commercially available TeraLight™ fiber for 2.5, 10 and 40Gbit/s transmission.

This breakthrough transmission would allow for the transport of 640,000 simultaneous high-bit rate ADSL Internet lines or over 100 million simultaneous voice calls on a single optical fiber. It will allow operators and service providers to dramatically increase in the future the capacity transported over a single fiber, thus further optimizing the cost of their networks.

To achieve such a performance, Alcatel has developed a new technique to optimally use the bandwidth of optical amplifiers. When carried over an optical fiber, wavelengths generate interferences. It is therefore necessary to have a minimum spacing between each wavelength in the networks, which limits the number of wavelengths that can be transported over the same fiber. The higher the bit rate, the stronger the interferences and the more limited the number of wavelengths. The innovating technique developed by Alcatel allows to minimize the effect of these interferences and to pack an unprecedented number of DWDM 40Gbit/s channels on a single fiber, thus reaching the world record capacity of 5.12Tbit/s. This laboratory demonstration also exploited Alcatel's leading-edge NZ-DSF optical fiber, TeraLight™, which offers an optimized design for DWDM systems. This result, achieved in realistic conditions with in-line amplifiers, underlines the Alcatel strategy to optimize all the key technologies required to offer cost-effective, advanced optical networks to cope with the Internet explosion.

---

## Sycamore's ODSI in Major Setback

**[From :Light Reading]** DENVER -- Efforts by Sycamore Networks Inc. (Nasdaq: SCMR) to gather industry support for its way of automating tomorrow's telecom networks has suffered a significant setback.

Its way of doing things was rejected by other vendors at an August 17th meeting of the Optical Internetworking Forum (OIF). And that's dented Sycamore's prospects of being first to market with standards-based software that enables edge equipment like routers to set up and tear down connections on demand over optical backbones.

It might also mark the beginning of the end for the Optical Domain Service Interconnect coalition (ODSI), a Sycamore-led group of vendors that was set up earlier this year to try and speed up the development of standards in this area. It was ODSI's proposals that got rejected at the recent OIF meeting, and that setback has led to some members questioning whether it's worth carrying on.

News of these developments have only materialized in the past day or so, when the OIF got around to publishing a press release about its meeting, which was held in Barcelona.

For the full story, please go to:  
[http://www.lightreading.com/document.asp?doc\\_id54](http://www.lightreading.com/document.asp?doc_id54)

# Gigabit Ethernet Empowers Fiber to Home

## Increasing Availability of Fiber from Utilities & CATV Enables Creation of Micro Telecommunications Providers -- World Wide Packets' Gigabit Ethernet Gear to Enable Delivery of Vast Bandwidth to Home and Small Business

**Editor's Note:** Bernard Daines is the founder, President and Chief Executive Officer of World Wide Packets. He is widely recognized as an expert in Ethernet technology. Daines founded World Wide Packets in 1999 to develop the first cost effective solution for broadband services for residential and commercial use. In 1994, Daines founded Packet Engines to develop leading edge Gigabit Ethernet products. Packet Engines was sold to Paris-based Alcatel in 1998 for \$325 million. In 1992, Daines cofounded Grand Junction Networks, a Fast Ethernet switch manufacturer with \$4 million in capital. That company was sold to Cisco Systems in 1995 for \$350 million. Prior to Grand Junction, Daines worked for IBM and Hewlett-Packard, and founded his first company, Tidewater Associates, which provided contract engineering services to many of the largest technology companies. Daines has created 38 successful ASIC designs, as well as numerous circuit boards, systems and test equipment products, including the first Ethernet products for 3Com in the early 1980's, and the first Ethernet LAN switch in 1988. We interviewed him on August 31, 2000.

**COOK Report:** With gigabit and soon ten-gigabit Ethernet coming on like gang busters at the periphery of the Internet, traffic loads on backbones seem poised to grow by several orders of magnitude. What is this doing to the future viability of backbones?

**Daines:** We believe that the core of the network is simplified by becoming all packet-based rather than the other technologies that are more constraining and difficult to manage. Consequently, as all of the deployed fiber gets lit up with wave division multiplexing, the capacity of the core will continue to grow. The whole notion of global WANs or Ethernet from end-to-end is somewhat counter to the prevailing thought from some years ago, that the WAN was error prone, temperamental, expensive, and took the high priesthood to operate. It is quite hard to get the telcos away from this outlook because this is where their bread and butter is: to charge high prices and not to innovate.

**COOK Report:** For the time being you are likely to have a lot of SONET in those backbones but do you also see some of them turning to Ethernet?

**Daines:** Absolutely. I see alternate service providers starting to light new fibers with Ethernet all the way.

**COOK Report:** Bill St Arnaud says that he sees three business models for backbones. The first is traditional telco circuit switching (in colors according to Sycamore networks). The second is MPLS where you are doing your routing at layer two rather than layer three. In the third is the optical border gateway protocol extension where users at the edge are responsible for establishing switching cross connects without the intervention of a carrier cloud in between.

**Daines:** The last is an attractive idea. Think about how FedEx routes packages. The post office used to sort packages multiple times along their routes while FedEx goes to great trouble to get complete routing data on pickup and sort only one time at Memphis.

**COOK Report:** This makes good sense. Now you mentioned a moment ago that people on the periphery with gigabit networks would build their own backbones. Would you describe what you have in mind?

### Getting Gigabit Ethernet Data to the Rest of the Internet

**Daines:** What I meant was that if you look at Qwest, Level 3, and Enron and those people who have fiber all over the place, it gets relatively easy to say, if we will light up someone else's strands we will light them with packets. Now for a while they may be lit with packet over SONET. When the 10 gigabit Ethernet standard is done, ten-gigabit Ethernet frames will be the technology of choice.

**COOK Report:** How are these business models to compete with each other? Gigabit and ten gigabit Ethernet is certainly less expensive than SONET. But there is hundreds of billions of dollars of SONET infrastructure in place worldwide. Most of the legacy Internet backbones are SONET. What happens when traffic needs to go from a new Ethernet backbone to an older SONET one?

**Daines:** We already have this kind of an issue with things that are ATM related and

things that aren't. It can be done at exchange points with Ethernet to SONET conversion.

**COOK Report:** What I meant was more how they would handle peering issues? How much will the gigabit edge of the Internet have to pay the older backbones for transit of their packets?

**Daines:** In a world like that you will measure your data very carefully and make routing decisions to keep traffic that an upstream transit network has to deliver to an absolute minimum. The criteria for switching points is going to be determined by the type of data and the routes available to it and by what the costs are at each exchange point. This is when you begin to see some advantage with software in the Junipers, Sycamores and Avicis, which enable decisions to be made within the packets as they go along.

**COOK Report:** Are you saying that, as the gigabit wide area networks get built out, and the MANs begin to link to each other, you might be able to get most of your traffic to UUNET customers in Philadelphia delivered by going on an Ethernet backbone to an exchange point in Philadelphia? Then at the exchange point you do a local handoff to UUNET's customers and don't need to worry about paying UUNET for any transit charges?

**Daines:** That is correct. But in the meantime, if you have packets for delivery to south Florida and this is an area that you don't yet cover, you might have to go to some switching point where you convert your packets from frames into cells and then make your recipient reconvert them at the other end. This kind of thing will go on for a while.

**COOK Report:** But if you are going to be doing Internet services, in theory you have to say I will deliver your mail any where. Consequently there may be some period of time where you will buy traffic from the old legacy backbones?

**Daines:** Sure. You may have to convert it to cells in order to be able to deliver it.

**COOK Report:** When you say that, is this tantamount to saying that you are putting it on their backbone and letting them carry it?

**Daines:** Yes. Now let's go back to the days when Sprint was advertising the first all digital network with the "so quiet you can hear a pin drop theme." Think about the same thing back in those days and note that this still happens in voice. If you don't happen to go there, they would at least circuit switch you and if a call went across something other than an all layer two network, there might be conversions from analogue to digital and back and forth two or three times. But generally this doesn't happen much anymore.

**COOK Report:** Are you aware of the extent to which major corporate networks are dropping their managed SONET services and going to dark fiber over Ethernet backbones? Are we getting as a result a huge amount of corporate data traveling independent of the Internet? Can this data be aggregated in such a way as to be able to be used to meet the peering requirements of the old legacy backbones?

**Daines:** Clearly, in answer to your first question, corporate unmanaged, point to point Ethernet traffic is clearly growing rapidly and I think will continue to do so. If you think about SONET, you will understand that SONET was put in there with all its bells and whistles to carry dumb data — namely digitized voice. You have a lot of overhead necessary to manage that dumb data. Now Ethernet is smart data and therefore you don't need all the complicated and expensive SONET infrastructure to carry it. These qualities too will, of course, be able to be carried forward.

Now with respect to peering, whatever qualifications someone like UUNET might have might seem formidable now, but if this corporate traffic continues to grow, and works like it does now, UUNET will have to begin to acknowledge reality and bring their peering agreements in line with these changed conditions.

**COOK Report:** Where do the exchange points of the kind of Equinix is building come into play? For example if you take Atlanta Georgia as a large regional hub might you have a business there. You might aggregate traffic from corporate backbones and traffic that your own company generates and then hand it off to the backbones of UUNet Sprint, AT&T, Genuity and so on?

**Daines:** For a period of time you will have to do things like that. But eventually what you will do is buy on a spot basis, or a short term basis, bandwidth between various points. Think about Enron, for example, getting into the brokerage of bandwidth. I think it will be essentially a bit like a toll highway where when you need to get packets from here to there you can choose to pay someone a small amount for their delivery.

## Managing Your Interconnection Topology

**COOK Report:** Suppose I'm a company in Atlanta Georgia during the summer of 2000, a Kmart, Wal-Mart, or a tire company and I'm getting my Internet connectivity from one of the older managed SONET networks to get traffic to me from these new Ethernet networks. Don't you have really only two choices? One that Ethernet traffic will have to go into a paid for transit link to one of the large legacy SONET backbones. Or, two, do you begin to see any signs of companies like this realizing that it's going economically advantageous for them to run their local loops to the nearest large exchange point in order to gain the efficiency of transport on Ethernet backbones?

**Daines:** If you think about the situation with Wal-Mart today, you will recall that it has all sorts of circuits with its customers that are used for Electronic Data Interchange and that these circuits are largely PVCs of one sort or another within the carrier's networks. Whether it's a color or T1 channel or a DSL whatever or an OC whatever channel they are PVC issues. What you are saying is that, at some point, it's going to be economical, instead of paying hundreds of thousands of dollars a month for the privacy and the absolute guarantee of the availability of that circuit with respect to the users, to pay a lesser amount for a predicted statistical average use of the Internet packet highway from point to point. And once that transition starts, they can be connected to other people in other places — as well as to the more open Internet at other locations. I think it's a gradual process. I think it will start when some big outfit says: OK we have full Ethernet from here to there and you can take branches off of that and expand and go parallel a month later.

**COOK Report:** when you talk about major Ethernet packet highways between cities, you are obviously talking about backbones. Are companies beginning to build these and if so who?

**Daines:** Sprint and others have said they are moving away from ATM to all packet. Now that is not exactly the same thing as moving to Ethernet native, but that is the first step. What we think is facilitating this, is the colors or lambda's obtainable from DWDM. If you think of what was necessary for a packet to get from one side of the country to the other, you will recall that until very recently, it was necessary to go from light back into electronics in order to do your reforming, reshaping, and error recovery or whatever and then back to light again. That would mean that each of those points would have

to understand what it was interpreting and restructuring.

Now, as you are able to take DWDM and amplify colors, the number of intelligent relay points goes down. In that transition state, you can go to the intermediate relay points by just amplifying the colors. With only the major relay points having to understand what it is, you have the opportunity to go in one circuit from Point A to Point B with color application in between. As you go to a point where we have to do electronic amplification, you bring out the various colors. Some of it can be reformatted and put on ATM SONET and the rest reformatted and put on Ethernet. And you have the opportunity to introduce the Ethernet or the packet over SONET and fewer and fewer points as you are building out the infrastructure.

**COOK Report:** I think I hear you saying is that the pendulum is probably swinging back from the points of private peering interconnection and going instead in the direction of public, modern and neutral fiber interconnects all along the lines of what Equinix doing.

**Daines:** Correct.

**COOK Report:** If you look at Williams, Level 3 and Qwest (which seems to be more like a telco than the first two) can we anticipate that some of them will begin to open Gigabit Ethernet links between their major backbone nodes? Or do you think they will let people come to them and buy pieces of fiber to light up with Gigabit Ethernet themselves?

**Daines:** I think that at first it will be the latter until the large fiber owners understand that the demand is such that they should be doing it themselves. You are already seeing it happen. Companies like 3rdWire, in Wayne, Pennsylvania <<http://www.3rdwire.com>> are popping up and building out from 30 to 50 cities with gigabit Ethernet to the end user and tying it together with fibers traded from other people. Telseon is another company that is building out fiber to the home in Palo Alto <<http://www.telseon.com>>.

The telco model had been to put fiber in for the last mile for \$50,000 and then spend \$500,000 on SONET gear to make it work. Today the model is to put fiber in for \$55,000 because prices have gone up but I need to spend only \$10,000 more on Gigabit Ethernet gear and I find that I have a viable operational data network. Suddenly this becomes very cost-effective. 3rdWire <<http://www.3rdwire.com>> has a viable business model, as does Yipes, by taking enterprise Ethernet gear and connecting local businesses. <<http://www.yipes.com>>

**COOK Report:** To what extent is the ocean going to be a barrier to this broadband expansion?

**Daines:** People are laying new transoceanic cables all the time. I think the bandwidth is there and as more and more traffic shows up at the shore, I think you will see some of the new fibers being converted to Ethernet. Alcatel which has over third of the under-sea cable market is focusing on dense wave division multiplexing of 40 gigabit lambdas.

**COOK Report:** Do you foresee any problems in connecting high-speed SONET fiber lengths to Gigabit Ethernet links?

## Enter Cable TV and the Utilities

**Daines:** It is working fairly smoothly. I think you are going to see the landscape is changing rapidly. One of the things that held down the spread of the cable TV market was the cost to the head end. The thing that is changed this equation is the use of hybrid fiber coax that saves you from having to go through many analog amplifiers. With the use of fiber you can branch out from the head end for up to 20 miles and therefore, share head ends between small communities. An estimated 50 million km of fiber being produced next year is bringing almost everything to within fiber's reach. When a hybrid fiber coax cable TV system is built out, it invariably includes several extra pairs of fiber. As a result, it becomes very easy to take fiber switching points right to the end user's neighborhood.

**COOK Report:** exchange point builders are multiplying. In addition to Equinix you have IBM and Qwest. Do you know how many companies are in this market and will a hybrid fiber coax solution gravitate toward an exchange point or can you use it independently of one?

**Daines:** Hybrid fiber coax is restricted entirely to the cable TV market. I'm just making the illustration of fiber invading and enabling the expansion of the old analog cable TV system. In many many communities fiber is getting a lot closer to homes than it ever has been before. This has become a starting point for doing what we're doing at World Wide Packets, which is taking fiber right to the very edge of the network.

**COOK Report:** So are you saying then that you not only have to be aware of the amount of carrier fiber but that you also have to be aware of the rather enormous amounts of fiber laid by the cable industry that in effect develops a parallel network to that of the carrier network connecting many cities ?

**Daines:** Yes and you also need to look at

outfits like Charter Communications which is a cable overbuilder that is putting in a second cable system on top of the one belonging to the incumbents. < <http://www.chartercom.com/gateway.asp>>. Other companies are Clear Channel <<http://www.clearchannel.com/calendar/cccbkgnd.htm>> and WOW <<http://www.wow-com.com/index.cfm>>. Also don't forget RCN on the east coast <<http://www.rcn.com/>>.

**COOK Report:** As we showed with an interview about VDN in Montreal last month, if you have a fiber infrastructure for the delivery of cable-TV, it becomes very attractive to put a Ethernet data network on top of that fiber.

**Daines:** Indeed. And another thing of which you may not be aware of is that by federal law the utility companies cannot use the communications circuits of the telephone companies to run their systems. They have to have their own communications circuits with which to tie together their generating and transport infrastructure. This has led to the utilities fiberizing their own rights of way. In the northwest United States, a lot of the utilities are in what is called Public utility districts (PUDs). These districts are usually contiguous with counties. Now 27 of these counties of gotten together to form a system called and NoaNet (Northwest Open Access Network) <<http://www.noanet.net>>.

We are seeing fiber pairs from Bonneville high tension lines linking these counties together. Once linked, the counties become an interesting buying block or a selling block, if you wish, for services. One of them, Grant County PUD, in Ephrata Wash., which is also one of the first ones that has our equipment in it, has adopted a business plan that will make it a carrier of carriers. They are not in the service business per say. What they will do instead is provide a fiber path to every building in the county during the next two years. Whoever wants to sell dialtone to the user there can do so. Whoever wants to sell Internet service can sell Internet service. And when you get away from the monopolistic set top boxes on the TV, whoever wants to sell entertainment on the fiber can sell entertainment. And by the way the PUD will read your electric meters is well.

Now some of these rural counties are under served in cable TV and they're even under served in telephones. Since deregulation when you build a house off the beaten path, unless you pay a huge amount of money, the phone company will not bring you phone service. But the utilities have now realized that that if they get electricity to you they can get fiber to you.

**COOK Report:** So a municipality with the right type of leadership can float a bond is-

sue to install a fiber network in the municipality. And as a result the municipality may run its own telecommunications operations and in doing so gain revenue from homes and businesses which can be plowed back into the needs of the community?

**Daines:** Absolutely. In the Grant County pilot project, known as ZIPP, they will provide nearly 100 residential test sites with Gigabit Ethernet broadband connectivity over their fiber optic network. They are currently using the World Wide Packets residential gateway unit to support voice, data and video services over Gigabit Ethernet links.

Also, if you think about the power companies, many of them are getting rid of the dams and the generators and are becoming merely carriers of power off of the grid. With deregulation, you as a consumer, can decide from which source you will buy your power, even though it all comes in over the same wires. In this new environment, with many utilities needing to find good use for their trucks and personnel, it leads to a changed business model that makes them quite willing to begin to install fiber.

**COOK Report:** Consequently in many communities in the U.S. you have not only carrier fiber and cable TV fiber but also utility fiber. In a given community is it relatively easy to find out what fiber runs where?

**Daines:** I think it's relatively easy but as the strands proliferate, it probably will get more difficult to keep track of them.

**COOK Report:** So in many places it is gone far easier than one might have assumed to bring fiber to the home. How then does this fit in with your business model? What exactly are you doing?

## World Wide Packets Business Model

**Daines:** We are making the equipment that can be used to bring gigabit IP into homes via fiber. The first device that we're making is the subscriber distribution unit (SDU) that brings in a pair of fibers and has all the features, bells, and whistles necessary to handle integrated video voice and data correctly. These are the units being installed in Grant County. The SDU has ports on it for a number of 10/100 Ethernet devices you might have as a subscriber, be they computers or set top boxes. Now to make the transition easier, you can also get it with a number of POTS lines so you can do something other than voice-over IP phones, although you can do those too. We can even aggregate those POTs lines and deliver them to the local telco in T1 fashion.

The fibers go from the subscriber distribution unit that would be in your home, or your hotel room, or your small office, to an aggregator box that would handle a couple of hundred subscribers, what we call the Community Distribution Unit (CDU)

**COOK Report:** Do you have any of the subscriber distribution unit boxes priced out yet?

**Daines:** Yes and no. Typically we don't sell the box independently. Although we do in the early days. With a complete system what you do is sell a link. By link we mean a connection from a port on the Community Distribution Unit that feeds all ports on the Subscriber Distribution Unit. We're priced a bit higher than hybrid fiber coax today but we're going to get very competitive with them very quickly. If you were offered a hybrid fiber coax system versus an Ethernet system, and you put all the costs together you would find out that there won't be much of a premium at all in a couple of years.

**COOK Report:** So in other words here in New Jersey where Comcast is the local cable provider, are you saying that in a year or two World Wide Packets might have an office in Newark and Trenton and Philadelphia and that you would have made your own arrangement with Comcast to ride on their fiber? It sounds like you're saying that you're going to go in and do something in parallel with these earlier technologies?

**Daines:** I think our customers will. We are predominantly hardware builders with a software system that goes with our hardware. We think our customers will be the utilities and the cable companies.

**COOK Report:** The competitive cable companies and people like RCN for example?

**Daines:** Yes. And of course once their competitors use our technology, the existing cable companies will have to fight back with something like it. Our other major customer is found among the builders of large housing developments. When you are building a tract of some several hundred homes it is very cost-effective to put fiber to the homes in the trenches that supply the houses with other utilities.

**COOK Report:** Does anyone have statistics on who these builders are and where they are building and how many houses they're putting in?

**Daines:** Some of the leaders are The Broadband Group <<http://www.broadbandgroup.com>>; Tylite <<http://www.tylite.com>>; and Vialight <<http://www.vialight.com>>. I would also refer you to Sweden where B2 is doing it with apartments and to the Benelux countries as well, where they have plans for several million

homes. Pacific Power in the state of Nevada has been very public about their plans to fiber homes. And then you have WIN (the Western Independent Network) which has raised a million dollars to take a gigabit fiber to residences. They have certain cities targeted for service.

**COOK Report:** Is what you are developing then primarily the box that would terminate these gigabit services in the home and small office? As well as another box that aggregates these services at something like a headend.

**Daines:** Yes. What you're calling a headend we are calling neighborhood box which will be meshed together in rings and then tied into a backbone — perhaps to someone's large router, or into someone's NOC, or at an exchange point. Different people will do this differently.

**COOK Report:** What is the thought on how this will be priced?

**Daines:** We expect to get cost of down to our utility customers to about \$1,200 a link. They could sell it to the user for a monthly fee or a one time charge. The business model will be up to them.

**COOK Report:** In other words if I'm a housing developer and I have a hundred and 20 houses that I want to connect into a neighborhood hub in a hub and spoke architecture, the total cost for the hardware would be about \$144,000 (120 times 1200). This would include the cost of the central hub and a subscriber distribution unit in each house but not the cost the fiber.

**Daines:** That's correct.

**COOK Report:** Wouldn't one business model be for me to offer this kind of connectivity to home owners for say \$5 less than the cable TV provider charges. When I do this I will explain to the owners that in connecting with me they would get a gigabit of data service essentially for free.

**Daines:** There are several startup businesses that have adopted these tactics. One of them is located in Issaquah, Washington. It is called ViaLight. This company is currently working with developers to bring dark fiber into new housing units in Washington state.

## Micro Telecommunications Utility Business Model

**COOK Report:** Is there an economic model of pricing for entertainment and telecommunications services that makes a fiber to the home feasible? From an economic point of

view does anything serious change when you get to an average household spending from about \$100 to \$120 a month on telephone, cable TV, cellular, and Internet? If we have people paying this much money for separate and non integrated services, does the business model now become one of offering them a set of integrated services for a similar price at a gigabit of bandwidth?

**Daines:** Sure. They might even offer it at a lower cost and limit the bandwidth for a while. And then of course just like pay-per-view or whatever, once you have that bandwidth in there, you have the consumer hooked up for a modest fee for basic services. Then you can get the incremental revenue for HDTV and video on demand, or high-speed Internet.

**COOK Report:** What you are enabling with this kind of business model and this technology is the creation of a bunch of micro telecommunication utilities.

**Daines:** Exactly.

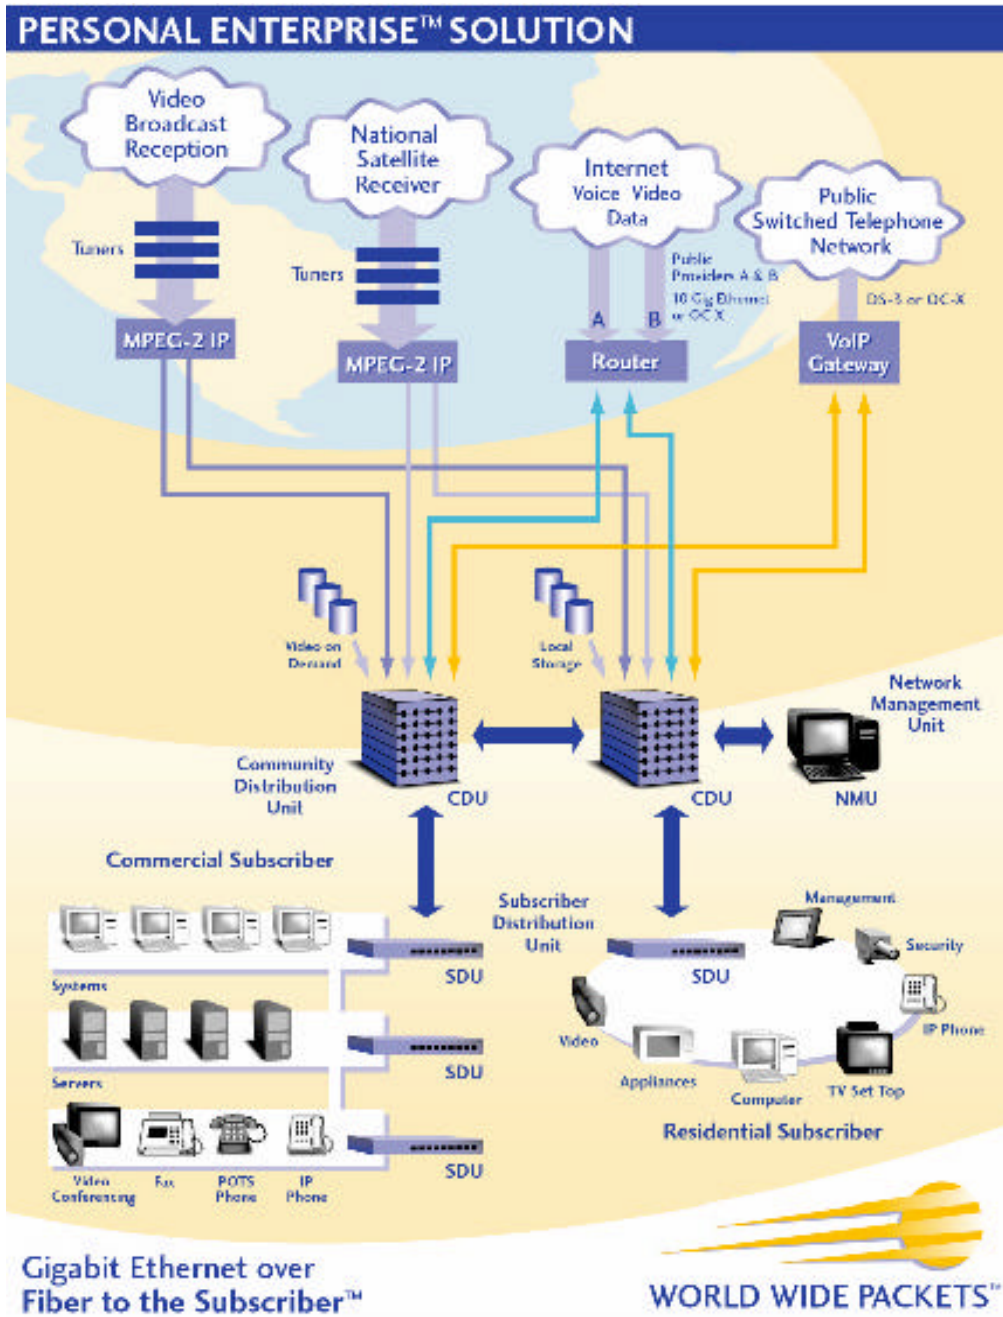
**COOK Report:** Please describe in greater detail the boxes you are building. Do they use primarily off-the-shelf components? From an engineering, manufacturing, distribution point of view what you doing? What are your challenges?

**Daines:** Our challenges are to build enough boxes for next year. But these boxes are not quite as garden variety as you might think they are. You can always use enterprise gear for this, but you're going to pay too much and the gear may not have the right features to do the integrated services that you want. We have created some very unique architectures to make all this work out and to have the cost come out where we need it to. If you were to go to the Foundrys and Extremes and other companies who are selling Gigabit Ethernet Switches, they want \$3,000 to \$5,000 more per port than Cisco. Our port cost on our aggregator box is going to be down in the five to \$600 per port range. So we are obviously doing something very different to cut costs by this much in order to enable the mass market.

**COOK Report:** The technologies are generally known. Therefore what you have to do is to select the right compounds and design and engineer them to work with each other in the least expensive way possible?

**Daines:** Even more than that. We are building our own ASICs to do the job correctly. We are in effect designing our own circuits and our claim to fame is that our engineering group has done dozens of custom circuit designs (ASICs) over the years.

**COOK Report:** Do you have a product that is actually finished?



**COOK Report:** So what you're saying is that if one looks at some of the fancier Cisco and Nortel Ethernet switches you are re-engineering less expensive and more robust versions of these to serve as neighborhood aggregators in your Personal Enterprise gigabit Ethernet networks?

**Daines:** Our goal is similar but we haven't reengineered anything on our customer distribution unit, rather we have it designed it from the ground up with an architecture that does the right job. The changes are primarily to the internal architecture. Bear in mind that I have built, during the last decade, a lot of the Ethernet boxes that Cisco sells. Consequently I've been through this before.

**COOK Report:** What do you see as the challenges that must be met in the next six months to a year in order to ensure the success of your enterprise?

**Daines:** We must execute well, our plans to build our boxes and equally important, we must do a good job of teaching our customers, many of whom have not been in the telecommunications business, to be able to deploy and operate this equipment in order to give their customers satisfactory service. Finally we must build and deliver sufficient boxes to meet customer demand.

**COOK Report:** It sounds like there is a role here for someone to be educating some of the utility people and home development builders as to what is now possible and can therefore become a viable part of their own business plans. Is this already going on?

**Daines:** We have had people out there doing pre marketing and pre sales for the last 12 months. Going to association meetings. Talking to utilities. Talking to builders. We will continue to do this.

## Afterword

On September 12, 2000 the Canarie News mail list offered some excellent pointers to these last mile issues.

The Computer Science and Telecommunications Board (CSTB) is carrying out a "Broadband Last Mile" study. There are some excellent papers on their Web site on different models for building community broadband networks and broadband to the home including fiber to the home.

**Daines:** We are beta testing our Subscriber Distribution Unit in that first public utility that I mentioned. And we have other equipment ready to go that is being tested by other groups.

**COOK Report:** What is the purpose of your network management software shown in the block diagram of your network on the opposite page?

**Daines:** The Network Management Unit (NMU) is basically, a network management software application. Fiber owners or network administrators would use the NMU to provision services to their subscribers, manage network operations, and define service levels. It provides a GUI interface, which can be Java-based or Web-based, to control the CDU and

the SDU—complete end-to-end management of their entire network.

**COOK Report:** To be viable do you also have to have your aggregation unit (Community Distribution Unit) released?

**Daines:** Our version of that won't be out of until the end of this year or the first of next year. What we are doing is to make a deal with one of the providers of OEM switching boxes. We can use a few of these boxes for demos and trials, although being much more expensive than what we will produce. They would not be cost-effective for installation and will not have all of our features. Nevertheless it does offer a way for us to get started.

# Deployment of IPv6 Held to Be Contrary to Interests of All But Largest ISPs

## Discussion Now Sees Deployment as Increasingly Likely But Because of Cost No One Wants to Be First

**Editor's Introduction:** Pressures to deploy IPv6 are mounting, driven in no small part by the desire to expand internet telephony through out the PSTN. As long as IPv4 and NAT boxes are predominant in the architecture of the Internet, the use of internet telephony applications will be restricted because the v4 internet no longer has anything resembling uniform end-to-end use of protocols. The debate however is complex and important because an IPv6 internet is potentially more friendly to centralized control by an entity tasked to maintain end-to-end protocol purity. We tend to be of the opinion the diversity of an IPv4 internet with NAT boxes is more desirable than an internet with uniform end to end IPv6 connectivity. A friend shared the following opinion. I am concerned that v6 (and it's end to end model with mapping to globally unique telephone numbers) would give the big networks an argument in favor of the continued dominance of their carrier cloud model. It would also offer the big carriers some financial advantages because of the cost to deploy IPv6 would be prohibitive for smaller ISPs. Think "PSTN" & "public Internet".... just like "public root servers" for the globally unique public address space." The article that follows summarizes some of the continuing debate on the deployability of IPv6.

On the IETF list on August 17th **Michael Richardson** observed: my service provider isn't IPv6 ready. I am doubtful that any of them will be for a long time as far as I can tell.

**Sean Doran:** Yes, a very long time, until at least the ambivalence in the IPv6 community towards established providers is resolved, and some large providers' and many small providers' concerns are dealt with.

Consider the rather nasty attitude in response to my technical deployment and utilization-scenario related questions raised here in the past 6 hours:

— likened to SPAM [Narten] — accused of trolling [Narten] — told to go away [Narten] — accused of not knowing the history of CIDR [Senie]

I could go on with other recent examples too, but you get the idea.

On the other hand, Itojun, like many others in the IPv6 community, answered politely

with useful information, and seem to be thinking about ways in which IPv6 could actually be used in practice, even with non-uniform support for native IPv6 by ISPs.

I assert that one of the critical stumbling blocks to acceptance of IPv6 by \*ANY\* large provider is the open hostility of many of the "leaders" of the IPNG effort towards those of us in the trenches.

Indeed, the whole ROAD process, and everything since has to me turned its back on [working with] people actually OPERATING the Internet as experienced resources who could help improve IPv6.

Perhaps if the IPNG leadership or at least their attitudes were changed, some useful engagement may happen.

Otherwise, Metzger's deployment scenario below is probably the only realistic one, because no business in its right mind would want to support a collection of people whose leaders openly accuse them of everything short of baby-killing.

Sean Doran <smd@ebone.net> (Working at a large ISP, incidentally)

**Perry Metzger** in <877liv1qq1.fsf@snark.piermont.com> on 03 Dec 1999 19:42:10 -0500:

We'll tunnel around you. You're irrelevant to deployment, anyway. [...] I gave up believing that providers would be paying close attention to the needs of customers about eight years ago [...] I'm certainly not going to hold my breath waiting for you guys to help deploying v6 now.

On August 22, **Thomas Narten** narten@raleigh.ibm.com Sean, Spamming the ietf list is bad form. Trolling is no more appropriate. Please take this elsewhere.

**Tim Salo:** Sean does have a habit of asking questions that highlight the fact that IPv6 isn't ready for wide-spread production deployment. I understand that you might you might not want to be reminded of this situation, but I believe that your response (particularly as an IETF area director) is inappropriate.

A more appropriate response might be to aggressively promote IPv4/IPv6 migration at IETF meetings. You might:

- o Coordinate an IPv6 migration help desk at the IETF that will help attendees upgrade their laptops to run IPv6,

- o Run IPv6 (only) on the desktop machines at the IETF,

- o Publish traffic statistics that compare the volume of IPv4 versus IPv6 usage at the IETF meetings,

- o Set an objective for when the IPv6 traffic is at least as great as IPv4 traffic at IETF meetings, and

- o Set an objective for when IETF meetings will support only IPv6.

If you can point me to a production-quality Windows 98 IPv6 stack, I would be happy to try to install it on my laptop, and maybe even run it at the next IETF meeting and help you with your migration project. (Oh, and make sure wireless works.)

Of course, I have been accused of being a counter-revolutionary for making these sorts of suggestions...

**[List member]** I think Trumpet have that .. trumpet.com.au Trumpet Winsock v5.0 is advertised to include IPv6 support.

**Salo:** In a bust of patriotism, I dutifully downloaded Winsock v5.0 and installed it on my Windows 98 laptop.

IPv4 stopped working. Initially, I assumed that this was a migration feature installed by IPv6 advocates. Upon further investigation, I learned that additional DHCP functionality is planned for future Winsock releases. Whatever that means, there doesn't seem to be enough DHCP functionality to allow me to connect to the network. (Winsock does, however, have a nice trace feature that let me know that the software at least called whatever BOOTP/DHCP code it does have.)

Not to be deterred by the complete lack of IPv4 connectivity, I plugged my laptop into an IPv6 cable at the NLANR/Internet2 Techs conference hosted by the University of Toronto. After a few more reboots, the Winsock software never seemed to have found an exit router. (While I am not entirely sure I believe the claims that IPv6 adds only minor changes to IPv4, I found IPv4 and IPv6 to be fairly equivalent in my brief

experience.) I will also admit that I don't really know if IPv6 is working. I think the four IPv6 cables have been carefully coiled in a pile for the whole conference. (Well, not quite. Just before me, somebody was hoping that the cables also supported IPv4 — apparently they don't.)

I uninstalled Winsock v5.0 and am thankful that I again have IPv4 connectivity, however politically incorrect.

So, if anyone else has any suggestions for getting IPv6 to run on my Windows 98 laptop let me know.

**Brian E Carpenter** IPv6 doesn't. IPv6 offers \*exactly\* the same QOS mechanisms as IPv4, namely IP Integrated Services and IP Differentiated Services. (There is also the flow label field in IPv6, but there are as yet no detailed specs of how it will be used and no false claims either).

**Vernon Schryver:** No false claims about IPv6 QoS? Absolutely none at all? No recent statements in this mailing list (or maybe it was end-to-end) that IPv6 QoS will be better than IPv4 QoS? No exaggerations in the trade press? Do you read the same trade rags and IETF lists I do?

Count the current applications of IP QoS in use by paying users. For that matter, count the successful large scale experiments. Recall what we have all been saying for the last decade. Add what the trade press has been saying, based on their honest (mis)understanding of progress in the IETF, vendors and on the Internet. Are you sure complaints about false advertising would be easy to deflect?

I hope and believe IP QoS will eventually be real and that the ATM stuff (including QoS) were empty promises for the IETF standard reasons, but as of today, the ATM guys have the high moral ground.

**Carpenter:** The doubling of the IPv6 address from 64 to 128 bits, was very specifically to enable an adequate (64 bit) locator component and an adequate (64 bit) identifier component in the address. And this was based on experience with several datagram network architectures of the past. The only realistic alternative was variable length addresses. But since we settled this in 1994, it seems somewhat beside the point.

**Schryver:** Yes, that's the spin I recall on the doubling of the IPng address. It wasn't an entirely dishonest gloss, but that's true of everything almost every committee does. I agree that the quick double-to-128-and-push-it-out-the-door-before-the-closed-questions-get-reopened was the least bad choice. I'm thankful that the base that was doubled wasn't only 64 and that it wasn't

more than doubled.

I don't think having a naked emperor is bad, provided his nudism doesn't force us to de-lude ourselves. However, political correctness and historical revisionism in the IETF is getting awfully thick.

To put it all another way, do you think IPv6 is on the schedule that was advertised 5-8 years ago, and if not, how much has it slipped? My recollection is that the advocates said "by 2000," the realists said "by 2003", and the rest of us said "by 2010 or 2015 at the earliest".

**Editor:** On September 1 the debate spilled over into the NANOG list.

**Nathan Lane:** The need for v6, at least for me, is in the deployment of mass quantities of end nodes. I do see a need for 4000 wireless end nodes at each of our sites - those could be serviced with a few app servers in the home office that knew v6 and v4-v6 tunnels [yes, technically these would be "routers"] at the remote sites.

At six to eight weeks for each router software release for the bug scrub and six months to roll out said code, we only have time to do it twice a year. Adding a new protocol would be formiddable unless tunnelling were used to virtualize the infrastructure.

**Roeland Meyer** 1) if one is using tunneling, then we really haven't gone over to IPv6. 2) if one is using tunneling, for leaf-nodes, then how is the core going to know those leaves are using IPv6?

3) Conservatively, every recent linux node has IPv6 kernel support and every Windoze box can do IPv6 (not to mention what we have just heard from the BSD camp <g>). This is over 50% of the leaves out there (conservatively). At what point is there sufficient market penetration of the technology to consider rolling IPv4/IPv6 interoperability/capability on the core routers and switches (something short of 100%, I would hope)?

4) Is it maybe that Sun, HP, Intel, Cisco, IBM, and the telco's, aren't all quite ready yet?

5) Admittedly, I haven't had the bandwidth to follow IPv6 all that closely. But, IPv6 has been pregnant quite a long time.

**Richard A. Steenbergen** There are two key problems which are preventing the widespread use of IPv6, IP Allocation, and network vendor support. Support for all hosts is actually one of the least of the problems.

Think of it this way, major networks want to deploy IPv6... How do they do it? They certainly can't do it on their primary back-

bone links and routers, the support from vendors is simply not there. Even if there was working code, they wouldn't dare deploy it on their production network, the code is too unstable (especially IPv6 routing protocols), and they risk looking unreliable in comparison to those who don't even make the attempt to support IPv6. So what do they do? Without the network there is not the demand for high traffic IPv6, and without the demand there is little desire to build the network. Should they buy separate routers, try desperately to make IPv6 work well on a spare 7200, and hope not to get a black eye from customers who expect the same level of routine-ness we are experienced with in IPv4? Should they provision more circuits because of this? Build a parallel network supporting IPv6, without a current customer demand? Or do they say, "we'll wait until the vendors get it right"? Even if you get past all that, there is still the very obvious fact that for certain vendors, the levels of performance we expect today are because of extensive tuning for IPv4 speed, and we know we won't be seeing this level of performance right off the bat from IPv6. And while they wait, there is no usable IPv6 infrastructure, less interest and development from vendors who don't really see a "need", and less demand from users who know there is no point when they can't find a network infrastructure to support them.

The other problem is IP allocation. ARIN and others are used to their comfortable role as IP Nazi ("No IP for you, come back, one year!"), they have happy extortion-scheme pricing on it, and they're not about to give it up voluntarily any more then NetSol would have without being forced... If you doubt this, go look at the current policies for obtaining IPv6 space.

**Sean Doran:** If I charge a customer more for IPv6 connectivity than for IPv4 connectivity, to offset the costs of dealing with ships-in-the-night routing (deploying it, training everyone to understand it), do you think my entire customer base is going to transition over to IPv6?

Ask yourself, as an ISP, how much more you are willing to pay your transit providers for IPv4 + IPv6 transit, and how you are going to get the money for that and for the deployment/retraining costs.

Then ask yourself, as an ISP, what benefit you get from IPv6. My answers: not a chance, none, and zero, respectively.

**Richard A Steenbergen** Now we're hitting the nail on the head. Well said Doran !!

**Leo Bicknell** Ask yourself, as an ISP if you can afford not to be testing IPv6 today. I don't think it's ready for deployment on any wide scale, due to a number of factors. However,

I do believe it will happen, and sooner rather than later, due to the demand for addresses. ISP's should not be offering it as a service yet, or charging extra for it, but they better be working on figuring out how it works, so when the day comes they can convert quickly.

**Nathan Lane:** Sean You speak in such extremes of your vision of what we should do with our applications and usage of IP. Might you expand your vision to include others?

How are you going to support a customer who legitimately needs a /8v4 worth of end-to-end connected devices and has the purchase orders to prove it? How would said customer, though, envision such an application if they couldn't get an allocation for it? They wouldn't and your business would never see it. The application would either die or find another provider or protocol. In my place, the application isn't dying. I've been fighting for SNMP proxies (to manage the existing plethora of devices), but these SNMP proxies don't exist and leave the wonderfully developed SQL databases so carefully designed over the years useless (and I hate proxies...I like end-to-end). Management doesn't like hearing you can't know the number of cans of Coca Cola in the machine outside store 359 at this "right now" second.

Read about Televend. A \$4.50 radio in vending machines. Ouch on the IP address usage for that application. How many vending machines do you cross day to day? We can NAT these to our hearts content, but eventually it must connect to the true supplier and that requires end-to-end. How can ipv4 support that load? I don't think it can.

Think ahead. My xeroxed copy of "RFC 1" is a fascinating journey as it was written shortly before my wife and I were born. It discussed the finances and realities of building the network from Santa Barbara and UCLA to Salt Lake City @ 1200 bps across the Mojave Desert. It discussed the costs of 2400 bps and the stations required to make it a reality.

Think about the future applications. My own three children do not know what a modem is. They expect and are delivered IP connectivity and if any of them reports network trouble, I know I have failed in my delivery of full connectivity for all applications. (Their access is filtered; but my son's very mention of DHCP sends my hackles rising.)

Business is now, indeed. Applications are later.

**Bill Fumerola** At some point, when ARIN just stops issuing IPv4 address space, I would say it would be in a ISPs best interest to switch. As an ASP, if my uplinks offered ipv6

address space/routing, I'd get it and start working with it `_now_` rather than later. — batz Have there been any studies done on IPv6 as an alternative to NAT? Besides IPsec, dynamic addressing, authentication and improved security, are there other benefits to deploying IPv6 instead of NAT?

On September 7 **Dana Hudes:** NAT is more than just a means to ease IP address space use. The use of a dynamic NAT pool allows the hiding of internal IP topology, thereby increasing security.

On Fri, 8 Sep 2000, **Masataka Ohta:** I'm assuming you're being facetious. Those features can be done with v4 and v6, using DHCP, IPsec, and ideally some other features.

If v6 is going to come into widespread use on the net, it has to be in production somewhere. Firewalled corporate networks are as sterile an environment to unleash it as any. As it stands, NAT was just a hack to conserve address space, and now that there are 'functional' v6 implementations maybe it's time to start thinking of a strategy for deployment.

The reason I was asking if a study was done was to find out if there was any good reason, beyond curiosity, to deploy v6 on private production networks. If not, how long should we expect to have to tolerate vendor hubris and bad hacks to get around depleted address space?

**Batz:** I\* (including but not limited to "T" and "IPv6") are facetious, of course. First, vendors of IPv6 address space should seriously tell vendors of Internet service supply IPv6 service.

Then, vendors of Internet service should seriously tell vendors of routers that they really supply IPv6 capable routers.

And there will be a v6-capable Internet, only after which there will be some good reason, beyond curiosity, to deploy v6 on private production networks. And then, we can get around depleted address space.

**Roeland M.J. Meyer:** This appears to be a classic chicken and egg issue. Which comes first? I am perfectly willing to deploy private-side IPv6, if I had a reason to do so. Without IPv6 support in the core, there seems little reason to do so. Perhaps, IPv6 substituting for NAT'd space? I don't know if it's even possible.

The bottom-line appears that everyone is waiting for everyone else to twitch first, then the shoot-out starts. However, no one is all that interested in twitching. It also appears that everyone seems to be pointing at the legacy /8's whenever the subject of IP allo-

cation shortages come up (with some possible justification). IPv6 seems to be a means of ignoring that problem and everyone knows it. The issue seems to be whether the consensus will allow us to ignore that problem and move on, or rat-hole on that problem while we live with IP rationing. The real question is whom is benefiting from sustaining the current situation?

**Doran** (referring to Meyer's statement: However, no one is all that interested in twitching.) Also, nobody is willing to get shot!

**Elliot Lear** Unless you're (a) a startup or (b) a VERY big company who can direct the market.

**Doran:** The deployment of IPv6 is going to be EXPENSIVE in terms of real opex and probably real capex as well, it IS going to be visible on the bottom line of every ISP on the planet, eroding whatever margins one has.

**Lear:** This is true.

**Doran:** I can't see the deployment of IPv6 \*ever\* leading to any but the shortest-term revenue upside (if even that), therefore until the entire aggregate gross revenue of transit-providing ISPs up and down the entire food-chain is threatened, nobody will be deploying v6.

**Lear:** This is false. In the end ISPs will be able to make a wash of it through pricing structures. First there are early adopters, and those are here now. That grows into a small group of networks. Those are likely to be here next year. All it takes are a handful of large ISPs to say, "I'm game", and it's amazing what the shape of the net looks like. Anyone who disputes this is disputing history, since this is precisely what happened with previous improvements, the invention of the FIXes, MAEs and BGP.

**Doran:** The only alternative scenario I can think of is the deployment of IPv6 by a large provider who believes it can trigger a huge consolidation by pushing smaller ISPs it is competing with into an expensive deployment through sheer hype.

**Lear:** You're half right. All it takes are a large providers with a real application.

**Doran:** I am inclined to believe that the second thinking is the REAL reason behind the recent announcements by a monopolistic and internationally expanding ISP in Asia that they will do an aggressive IPv6 deployment.

**Meyer:** The real question is whom is benefiting from sustaining the current situation?

**Doran:** Everyone who wants cheap, sustainable Internet transit, with the continuation

of plummeting prices combined with soaring available bandwidths. Introducing a whole new protocol requiring massive global operational changes is going to push up consumer prices and stall on investment in available bandwidths. There are only so many people and dollars out there, and one or the other is inevitably spread pretty thin in the current market.

**Lear:** I guess the major point your missing is that once a major provider goes the others are going to realize that the Internet isn't going to shrink in size, and so the cost of moving to v6 is only going to go up. Hit critical mass and the party really begins.

**Doan:** The scariest thing to an IPv6-Lover is that an early deployment is not to anyone's advantage because until there is real uptake by a sizable number of ISPs, the exact changes required on the dynamic routing side are simply not obvious, although the fact that the two protocols will run ships in the night is (e.g. www.microsoft.com works just fine with IPv6 but you see a blackhole with IPv4. ftp.cdrom.com's IPv6 path is much slower and lossier from where your customer is than ftp.cdrom.com's IPv4 path. Have fun finding and fixing the square of the number of problems you observe now, kids!

**Lear:** I don't see this happening. I see a lot of interest in running two separate networks, where they might get merged later.

**Doran:** In other words, it's all risk and absolutely no reward, and until it really honestly IS impossible to do hacks around the IPv4 shortage, nobody will deploy IPv6.

**Lear:** The other downside you fail to mention is the growth of customers. It's still positive, surprise. And when is the problem more manageable? A further downside is the cost of IP address administration that continues to climb. If the assigning authorities can allocate out larger blocks the cost of remaining at IPv4 becomes far more obvious.

**Doran:** Now, ironically, in the whole IPv6 selection process in the IETF, there were multiple proposals which paid a considerable amount of attention to the problems of partial, incremental deployment at the initial design level. CATNIP in particular was clever, because it provided not just a new packet format (which is all IPv6 did), but also a strategy to transition to practically ANY new packet format, should the initial assumptions about the pervasiveness of IPX and CLNP be wrong (which they were).

IPv6's initial assumptions are WRONG (we will die from routing dynamicism long before we die of IPv4 address depletion), and there is NO mechanism whatsoever to abandon the IPv6 packet format even at the primi-

tive level of curiosity-based micro deployment we see now.

Who will take the chance of a huge investment in managing IPv6 deployment, when it is not a given that IPv6 really will be the header networks will use after IPv4? We're talking about stranded assets being the only thing one gets for the money.

**Lear:** As it stands today there are 0 alternatives that are being seriously considered by all parties, so while it's not a sure bet it is a good bet to place some amount of time into and charge a premium for early adopters (remember those)? What do the early adopters get? A chance to reduce the number of times they will need to renumber, a task which today is still quite expensive, even with tools such as DHCP. Bill Manning's dream is still far from reality.

Also, while I believe that routing is a very serious scaling issue, the last time I looked, the growth rate in address allocation was picking up. This Is Bad <tm>.

This is not to say that I believe IPv6 to be the cat's meow. I think the proponents' marketing has been nearly as bad as your anti-marketing. It would be worse except as of late they seemed to have toned down where you haven't. I think the one area that will give people pause will be the header size. To me that's serious overhead for the often bespoken "killer app."

**Thomas Marshall Eubanks:** What about wireless IP? Isn't everyone supposed to be forced to adopt IPv6 once billions of mobile units start using it ?

<http://www.zdnet.com/intweek/stories/news/0,4164,2590226,00.html>

**Doran** (Sept 10) : So, the problem here is that the very smart doctors of the IETF observe a number of people complaining about hunger pangs, because they cannot get the number of ROUTABLE addresses that they want.

Enter IPv6, the Morphine of the Internet. The bright side is that, like morphine, it makes the patient much less hungry, but there are unpleasant side-effects, and no added nutrition. It also distracts one from real problems by introducing an euphoric "solution" to a migraine headache.

Sure, you can have all the addresses you can ever want, and more, but the fact is that they are not ROUTABLE addresses. A 128-bit number is no more useful than a 32-bit number if nobody can use it to contact you.

The unfortunate thing about IPv6-Morphine is that like any narcotic one sees things in a dreamy, rosy state, tinted with denial. Those

good doctors took away my pain, so they can take away any future pain too. Oh, but alas, the fatal disease IPv4 users are faced with, namely the failure of the global routing system to cope with increasingly complex and increasingly dynamic network topology, is not cured by IPv6-Morphine, it is EXACERBATED by it.

Like many narcotics junkies, the wireless folks are simply of the belief that the dealers are their friends.

**Jon Allen Boone:** What do you think the solution is? Or, at least, in which direction does it lie?

**Doran:** The solution is in a complete separation of "what" from "where", and using the "what" part to identify conversations INDEPENDENTLY from the topology across which they are desired/occurring. There are many ways to get there. More on that after dinner...

**Meyer:** I tend to agree here that routing is one of our largest bugga-boos. What we have is held together with spit, baling-wire, and liberal amounts of the "racer's edge" (duct-tape). With CERF.NET bouncing all over the place, for the past three weeks, maybe I've become hyper-sensitive to those issues. But, it appears to me that the entire BGP system is a very brittle patch.

It is for this reason that I recommend ABOVE.NET to all my globally visible portal/ASP/B2B clients. However, this doesn't relieve the problem of the end-user being outside the ABOVE.NET system and having to live with cold-potato routing, for uploading files to the site. We have too few public peering-points and they are under-sized (what happened to the regional NAP idea?).

What you are saying is that going to IPv6 is a one-way function?

I've actually looked at some of this. At the risk of ridicule, may I mention Flemming's IPv8? It nested IPv4 inside the packet, as a sub-set, and actually planned for co-existence and inter-operability with IPv4. It also answered a LOT of routing issues. Regardless of specific implementation, that seems like a more prudent approach. Does anyone know why such an approach-policy wasn't followed by the IPv6 team? Yes, I agree that CATNIP was also clever. With a little work, it "could have been a contender" (Brando<g>). The choice of IPv4 and IPv6 shouldn't be an XOR function and the point remains that transition was not a consideration of the IPv6 design (this is obvious). In most commercial shops, such an approach would not have been accepted/tolerated. Let alone, win any sort of design contest, as did IPv6.

# More ICANN Footnotes

Editor's Note: Dear Readers, after having spared you from immersion in ICANN's transgressions last month, we include a fresh set of ICANN footnotes. This material is primarily from Karl Auerbach and Ted Byfield and serves to document some of the issues behind the Member at Large Elections. Finally it concludes with the first part of Larry Lessig's campaign platform.

On August 6th Karl Auerbach made the following submission to ICANN's Independent Review panel. We include it because it shows very forcefully how ICANN ignores its own rules and obligations

## Subject: Submission to Independent Review Panel

This is a request for the ICANN's Independent Review Panel to review ICANN's rejection of Reconsider Request 99-4. This is [also] a submission to ICANN's Independent Review Panel under section 6.1 of ICANN's Independent Review Policy on the grounds that ICANN's "Board has acted or failed to act in a manner contrary to the Corporation's Articles of Incorporation and/or Bylaws." I have been materially affected by the contested action and I have exhausted ICANN's internal reconsideration process.

History: During the late hours of November 3, 1999 the final drafts of a package of agreements were worked out between Department of Commerce, Network Solutions Inc., and ICANN regarding the the Internet's Domain Name System. The next morning, November 4, 1999, in a meeting that began at 10am, the ICANN Board of Directors adopted this package of agreements. These agreements were signed on November 10, 1999. This package is visible on the ICANN web site at: <http://www.icann.org/nsi/nsi-agreements.htm>

Neither the drafts nor any of the final agreements were ever submitted to the ICANN Domain Name Supporting Organization (DNSO) for its consideration.

I submitted a timely Request for Reconsideration on November 17, 1999 complaining that the failure to involve the DNSO was a violation of ICANN's by-laws. The request was assigned number 99-4.

My request is visible on the ICANN web site at: <http://www.icann.org/reconsideration/auerbach-request-17nov99.htm>

Nearly six months later, on May 2, 2000 I received an e-mail from the Reconsideration Committee informing me of the committee's rejection of my request. This e-mail is visible on the

ICANN web site at: <http://www.icann.org/reconsideration/rc99-4.htm>

On May 2, 2000 the board adopted the Committee's rejection. The board's action is visible on the ICANN web site at: <http://www.icann.org/minutes/prelim-report-04may00.htm#00.30> It is unknown on what date the May 2, 2000 board meeting minutes were published.

### This Request for Independent Review:

The Reconsideration Committee ruled against my Reconsideration Request. That ruling should be overturned for the following reasons.

**First:** The committee relies on ICANN's Bylaw Article VI, Section 2(g) as a means by which the board may bypass the obligations imposed by Bylaw Article VI Sections 2(b), 2(c), 2(d), 2(e), and 2(f).

The committee argues that: This provision [Bylaw Article VI, Section 2(g)] confers on the Board the ability to act when necessary or appropriate to promote ICANN's purposes, even without a referral to one or more of the supporting organizations.

That argument is flawed. There is a well known rule that rejects interpretations of legal documents in a way that reduces portions of that document to mere surplus.

The committee's rationale reduces Article VI Sections 2(b), 2(c), 2(d), 2(e), and 2(f) to nothing but advisory language. The committee's rationale goes so far as to destroy the "primary responsibility" of ICANN's Supporting Organizations for matters within their respective spheres.

I submit to the Independent Review Panel that the committee's interpretation is improper. I submit that the proper interpretation is that Bylaw Article VI Sections 2(g) is merely an escape provision to cover situations not handled by the more specific rules established by Bylaw Article VI Sections 2(b), 2(c), 2(d), 2(e), and 2(f).

By-laws are rules for the corporation to follow. The ICANN board and the Independent Review Committee should not be permitted to reduce the By-laws to something that the ICANN board may follow or ignore at its whim. Speaking of whims let us move on...

**Second:** The Committee recites that the Board acted on the powers granted by Bylaw Article VI, Section 2(g).

Section 2(g) reads: Nothing in this Section 2 is intended to limit the powers of the Board or the Corporation to act on matters not within the scope of primary responsibility of a Supporting Organization or to take actions that the Board finds are necessary or appropriate to further the purposes of the Corporation.

This section releases the limits on Board powers in either of two situations: (1) - When handling "matters not within the scope of primary responsibility of a Supporting Organization" (2) - When

taking "actions that the Board finds are necessary or appropriate to further the purposes of the Corporation."

The committee did not deny that the Domain Name System and the agreements are within the "primary responsibility of a Supporting Organization". Thus the first of the above mentioned two situations did not occur. In fact, the committee's report clearly states that it relies on the second of those two situations.

If one reads Bylaw Article VI, Section 2(g) one will notice the words "that the Board finds are necessary or appropriate to further the purposes of the Corporation."

In other words, before using the second situation of Bylaw Article VI, Section 2(g), the board must make a finding that the proposed action is "necessary or appropriate to further the purposes of the Corporation."

ICANN's board has never made such a finding. The matter of making such a finding has never been on ICANN's agenda nor does it appear in any minutes. And as a witness to ICANN's Board meeting of November 4, 1999, I saw no discussion of this matter nor any unrecorded finding. In fact the first mention of the board's reliance on Bylaw Article VI, Section 2(g) appears nearly six months after the fact in the reconsideration committee's rejection.

**Third:** The Reconsideration Committee stated in its rejection that:

The Committee also notes that the ICANN/NSI/DoC agreements were posted for public review and comment for over a month prior to their approval.

This is neither relevant and nor factually correct. It is not relevant because it has no bearing whatsoever on the responsibility of the board to follow Bylaw Article VI Sections 2(b), 2(c), 2(d), 2(e), and 2(f). It is factually incorrect because the agreement was not worked out until the late night or early morning before they were adopted.

A period of at most from midnight to 10am - the period between the completion of the draft agreements and the start of the board meeting at which those agreements were adopted - does not even come close to the "over a month" that the committee claims. To the contrary, the public had virtually no ability during those ten hours to review the documents or make comment.

Fourth: Despite the reconsideration committee's assertion to the contrary, my contractual rights, present and future, have been changed by the suite of agreements entered into between ICANN, NSI, and NTIA. Those agreements impose many new conditions on domain name owners and their options for holding those names. To say that those agreements do not materially affect domain name owners, such as myself, is to trivialize the massive negative effect on stability of domain name ownership that has occurred as a result of these agreements.

The committee, to its credit, did proceed to the

merits of my complaint. And so should the Independent Review Panel.

## The ICANN Membership at Large Elections

Editor's Note: Karl Auerbach has done yeoman's labor in documenting ICANN's perversity, first in the election process itself and then with his platform for the members at large election. ICANN created the membership in such a way that only it would possess member registration data - vitiating in the process any hope that the membership could become a force independent of ICANN's leadership. As Auerbach points out in the excerpted correspondence with Louis Touton below trying to campaign for office without being able to identify and reach the electorate is a frustrating experience.

Date: Thu, 17 Aug 2000 16:24:32 -0700 (Pacific Daylight Time) From: Karl Auerbach <karl@cavebear.com> To: Louis Touton <touton@icann.org> Cc: edyson@edventure.com, Andrew McLaughlin <mclaughlin@pobox.com>, Michael Roberts <roberts@icann.org> Subject: Request based on being a candidate

I am a candidate running for a seat on ICANN's board of directors in the election from North America.

As a candidate, I request that ICANN make available to me the names and contact information (including e-mail addresses) of at-large members. I am requesting the list, in electronic form, of those at-large members in North America. (If necessary I am willing to cull those out of the full list if that would be easier for ICANN.) I would have no objection to the list being made available to other candidates at the same time.

I will use the list:

a) in conjunction with election to communicate with the electors in furtherance of my candidacy or

b) to understand the voting pool (in order to be more responsive to member needs) and to independently review the voter rolls for voter fraud, etc.

I would not retain a copy after the election had been completed and the results settled. Nor would I make the list available to anyone who was not a candidate.

ICANN's failure to provide these names and addresses in a timely manner will cause irreparable harm to my candidacy and to the election process. It will cause irreparable harm to my ability to interact with the electors for the purpose of discussing issues, formulating positions, evaluating campaign strategy, and obtaining both "member-nomination" endorsements and, later on, votes.

Date: Thu, 17 Aug 2000 16:26:32 -0700 (Pacific Daylight Time) From: Karl Auerbach

<karl@cavebear.com> To: Louis Touton <touton@icann.org> Cc: edyson@edventure.com, Andrew McLaughlin <mclaughlin@pobox.com>, Michael Roberts <roberts@icann.org> Subject: Request based on being a member

I am an ICANN "at-large" member living in North America. I have "activated" my membership. With respect to the now ongoing election process by which the at-large membership will vote to elect people to fill seats on the Board of Directors, I have the right under the ICANN bylaws and resolutions passed pursuant to those bylaws to submit a "member-nomination", make a "member-endorsement", and to cast a vote in the election..

I am a member of ICANN, as defined by California Corporations Code section 5056. I hereby request that ICANN make available to me for inspection and copying the names and addresses of ICANN's North American members, as required by California Corporations section 6330. In particular, I am giving notice under section 6330(a)(1).

I will use the list:

a) in conjunction with the nominations and election process to communicate with and interact with my fellow electors on matters pertaining to the nomination and election. Such matters include, but are not limited to, discussing issues, formulating issues to be presented to candidates, discussing candidate responses, and discussing voting strategy, or

b) to understand the voting pool (in order to be build communities of interest, formulate voting strategy, etc) and to independently review the voter rolls for voter fraud, etc.

ICANN's failure to provide these names and addresses in a timely manner will cause irreparable harm to my ability to fully exercise my electoral franchise in this election and will cause irreparable harm to the election process as a whole.

Mon, 21 Aug 2000 17:41:32 -0700 From: Louis Touton <touton@icann.org> To: Karl Auerbach <karl@CaveBear.com> Subject: Your Requests for Names and Contact Information

Dear Mr. Auerbach:

I am responding to your two e-mail inquiries dated 17 Aug 2000. Your inquiries request that you be provided with the name and contact information (including e-mail addresses) for voters who later this year will cast votes for nominees to ICANN's board of directors. Your first request asserts that you are entitled to this personal voter information by virtue of the fact that you are a candidate for nomination to the ICANN's board of directors. Your second request asserts that you are entitled to this information by virtue of the fact that you are an ICANN at-large "member."

The information you are seeking reveals personal details about tens of thousands of individuals

throughout the world who have signed up to participate in the ICANN at-large program. That program, which is designed to provide Internet users globally with at-large representation on ICANN's board of directors, was developed through a 19-month process with extensive input and participation from the Internet community. Throughout that process, concerns were regularly and forcefully expressed that individual voters in the program should not be subjected to unchecked waves of unsolicited e-mails appealing for support, particularly when any individual anywhere in the world is entitled to become a candidate. With over 158,000 at-large applications, and 161 candidates, the concerns about unsolicited e-mail appear to have been well-founded. As a result of these concerns, the at-large selection plan that was adopted provides for voter education about the candidates to occur through web sites provided by the candidates (and web pages provided for each candidate by ICANN), combined with periodic, brief, and neutral e-mailings sent by ICANN referring the voters to those web pages and sites. To preserve voter privacy, the selection plan specifically provides that candidates themselves should not have access to the at-large voter lists or email addresses.

As to your first request, I am not aware of any legal basis on which a candidate running for a seat on ICANN's board of directors is entitled to the personal information about voters that you are seeking. To the contrary, in response to concerns over personal privacy, the ICANN Board, Membership Advisory Committee, and Election Committee each has determined that ICANN should not facilitate candidates sending unsolicited e-mail to the voters, and ICANN has undertaken the responsibility to keep the members informed about the process and the location of the candidates' web pages, in lieu of giving candidates access to voter lists. You are aware of your ability to communicate with members in this fashion, and we encourage you to do so.

With regard to your second request, Article II, Section 1 of ICANN's bylaws states:

The Corporation shall not have members as defined in the California Nonprofit Public Benefit Corporation Law ("CNPBCL"), notwithstanding the use of the term "Member" in these bylaws, in a selection plan adopted by Board resolution, or in any other action of the Board. Instead, the Corporation shall allow individuals (described in these bylaws as "Members") to participate in the activities of the Corporation as described in this Article II and in a selection plan adopted by Board resolution, and only to the extent set forth in this Article II and in a selection plan adopted by Board resolution.

In view of this section, ICANN's "members" are nonstatutory ones, and the California Corporations Code provisions cited in your second request as supporting a statutory member's right to membership lists are inapplicable. I am not aware of any other legal basis on which an ICANN at-large "member" is entitled to receive personal information about the voters, and ICANN has not provided this information to any

other “members” or board candidates.

[Snip] Louis Touton Vice President, Secretary and General Counsel

Date: Thu, 31 Aug 2000 17:41:32 -0700 From: Louis Touton <touton@icann.org> To: **Karl Auerbach** <karl@CaveBear.com> Subject: Communicating with Voters

Dear Mr. Auerbach:

You did not respond to my e-mail of August 24, 2000, seeking clarification of your position. Therefore, I am writing in further response to your e-mails of Thursday, August 17 and Monday, August 21, 2000 without the benefit of your further input.

The At Large Director Selection Plan that has been adopted provides candidates with effective means to provide information to, and communicate with, the voters, without compromising the privacy of the voters or subjecting them to a deluge of unsolicited e-mails from the numerous candidates. I’m writing this to make sure you are aware of these alternative mechanisms.

1. One such currently active mechanism provides for education about each of the candidates through the ICANN-operated website at <http://members.icann.org/nom.html>, with a link to the North American candidates, including you, at <http://members.icann.org/nom/cp/NA.html>.

2. Each candidate can, as you have, provide a link to his or her own personal website to further educate and/or interact with the at-large voters. Thus, on your own website, you are free to solicit voter e-mail or other information from those visiting your site, set up message or discussion boards, set up scheduled chats, etc.

3. Throughout the endorsement period, ICANN is sending to all at-large voters (i.e. those who have completed the registration and activation processes) approximately once per week (for a total of five mailings) a neutral e-mail listing of the names of individuals seeking to be endorsed in their region. The e-mail contains a link pointing to a web page from which the voter can hyperlink to candidate web pages, as well as a link to the page from which on-line endorsements are made.

4. Throughout the voter-education period (from September 9 to the end of the voting), ICANN is sending all at-large voters approximately once per week (for a total of approximately 4-5 mailings) a neutral e-mail listing the nominees on the on-line ballot and providing links to their web pages.

5. Additionally, during the voter-education period, ICANN intends to provide a public question and answer forum for the candidates, during which voters may pose questions to the candidates. Each candidate will be the moderator of his/her own forum, choosing which questions to answer and how to answer them.

In short, the system ICANN has implemented is

appropriate and, in fact, is working as intended. Any voter wanting information about you or any other candidate has easy access to the above pages and receives a weekly e-mail reminder to review them. Candidates can use some or all of these means to communicate with each other and/or the voters regarding nomination and/or voting issues.

[Snip] Regards, Louis Touton Vice President, Secretary and General Counsel

Date: Thu, 31 Aug 2000 23:52:01 -0700 (Pacific Daylight Time)

**From: Karl Auerbach**

To: Louis Touton <touton@icann.org>

Subject: Re: Communicating with Voters

[Karl references Toutons e-mail of August 24, 2000, and his non reponse there to.]

I’m busy running a campaign and I do not have any indications that ICANN has any intention of responding in any way that conforms to the California statute that allows members to obtain the membership lists and use them for purposes associated with the corporation. The list of ICANN “alternatives” that you sent me makes it clear that my assessment was correct.

(I know that you assert that we “members” and not “members” but I disagree.)

The need for access to such membership lists is patently obvious - as it was to the committee that wrote the model statute and to the legislatures of several states, including California, that mandate such access.

Members need not disclose to management how they intend to use the membership lists except to say that such use will be in association with the corporation. And as you know, California gives the corporation clear authority to take action if a membership list is misused.

I reject any assertion that what you have described below constitutes a reasonable substitute, as contemplated by the California statute, for access to the membership list.

I may note also that even that list below isn’t being met by ICANN. For instance, even though the “member-endorsement period is nearly 70% complete, ICANN has sent only one fifth of the mailings that are promised. And there are still many people who still have not received their PINs.

But even if ICANN were to flawlessly perform all of the promises below they would not amount to anything but a management controlled house organ.

In addition, given the rather questionable technology used for the registration and endorsements, not to mention the fact that the software is rumored to have been written by one who is at best highly partisan and at worst one who has admitted to having created multiple electronic identities in his participation at ICANN, access to the membership list is critical to detect manipu-

lation or error.

But I don’t have to further justify the need - California says that we members can have the membership list - and ICANN can not supersede California statute. All that is necessary is that the use of the list be reasonably related to the corporation. And there is nothing as related to a corporations life than an election to fill seats on the corporation’s Board of Directors.

And as I said in my request, I intend to use the list in conjunction with my quest for “member-endorsements” and, since it appears that I’m going to be on the ballot, to obtain votes. More than five days have elapsed since I made my initial request for the membership list. And I have received no response that I consider to be even an approximation of a reasonable alternative.

## Karl Auerbach’s Election Platform

**Editor’s Note** — Auerbach’s Member at large election platform may be seen at: <http://www.cavebear.com/ialc/platform.htm> It is a well reasoned provocative document that does a very good job of compiling a comprehensive list of ICANN’s defects. Some excerpts follow.

### Overall point-of-view:

As a general matter, I value procedure very highly - It is my belief that good procedure leads to good decisions. Now, I don’t mean that a “good” decision is one that satisfies everyone. Rather, I believe that a “good” decision is one in which every person feels that he or she has had a fair chance to make their position heard, that the decision maker actually did consider all points of view without prejudice, and that those who didn’t get what they wanted are willing to accept the decision.

Thus, the larger part of my platform is the reformation of ICANN and its procedures. However, I also believe that ICANN will be unable to be reformed until certain personnel changes are made.

And I do have some opinions on policy matters themselves.

To a large extent I am a firm believer that government is the proper embodiment for public debate and public policymaking - to my mind a body such as ICANN is Internet government no matter how many “private corporation” flags it might try to fly.

And I am not one who reacts with an instant “way to go” when someone utters the words “deregulation” or “privatization”. My studies of our history have taught me that private industry and private behavior are not always for the general good, that regulation, indeed governmental regulation, must sometimes be imposed to protect the public from certain self-interested actions of the few.

On the other hand, I do believe that governmental bodies ought to keep their fingers out of things

that don't need the active oversight and regulation. To a large degree I feel that the Domain Name System of the Internet needs no regulation, that economic and social forces will cause it to evolve in "the right way" without the intrusive and controlling efforts of a governance body such as ICANN.

As such you may find that my point of view is bimodal - If I find that something needs regulation, I tend to believe that it ought to be done by a governmental body, or at least a body that isn't afraid to recognize that it is like a government (and thus must operate with notions of "due process".) But if I find that something does not presently need regulation, I tend to take a rather more libertarian, hands-off approach.

My own personal politics tend towards the liberal or "green" point of view.

But when it comes to financial matters, I firmly believe in sound practices with strict controls. I have found ICANN to be lacking in financial common sense.

Oh yes, one more thing - I am willing to be convinced that a position I might espouse is ill advised and I'm willing to change my mind.

#### **Reform of ICANN - Openness/transparency/accountability**

An "open" process is one in which all interested parties can participate in a meaningful way and as equals to all other parties.

A "transparent" process is one in which the entire decision making process, from inception to closure, is revealed and recorded. To be fully transparent, a process must reveal inputs, issues, criteria, biases, misunderstandings, evolution of decision maker positions, compromises, votes taken, etc, etc.

An accountable decision maker is one who is both identifiable and can be held to account for his/her decisions. Board members are typically made accountable by elections and recalls - assuming that the electorate can obtain enough information to evaluate how their board members have performed. Staff members must be held accountable by the board.

One could look long and hard. But one is unlikely to find an organization that spends as much effort as ICANN does actively rejecting these principles.

It is my position that ICANN must operate with absolute openness, transparency, and accountability. This means that absolutely every input, every discussion, every decision - everything - must be done in open session with a written or electronic record. All decisions must be made by recorded vote - with the position of each director clearly shown.

The only exception to this would be matters pertaining to personnel and litigation. And even a decision that a matter falls into those categories must be made in public.

In addition, actions by "staff" are the epitome of non transparent and non accountable decision making. ICANN's staff must be required to operate according to the same principles and the board must be obligated to adopt, on the record, staff actions.

#### **Reform of ICANN - Jones Day must go**

Jones, Day, Reavis & Pogue is ICANN's law firm, and has been so since the day of ICANN's birth. Indeed Jones-Day actually performed the incorporation ceremony in its Los Angeles offices.

Jones, Day, in the person of its principle man-on-the-ICANN-scene, Joe Sims, was present for at least half a year before ICANN was born, working in the shadows, responding to unknown interests and possibly making unknown deals. About all we know about that period is that those who were not insiders to Joe Sims process were ignored and that those who objected were treated with condescension and abuse.

Over the life of ICANN, Jones, Day has been the the dominant creditor of ICANN.

Even now Jones, Day continues to receive a lion's share of every dollar that flows into ICANN. And one of Jones, Day's partners, Louis Touton, left the firm to become ICANN's Vice-President, Secretary, and General Counsel. There is in my mind a question about the appearance of propriety.

As a member of the board of directors I would have the right to rely upon the expertise of entities such as ICANN's law firm - that is unless I felt that I was not receiving the degree and quality of advice that, in my opinion, I felt would be necessary for me to properly exercise my duties. Given my interactions with Jones Day, I do not feel comfortable relying on their work.

There are many good people at Jones, Day, and I have no doubt that much, if not all, of its work is adequate. But the firm has no special credentials to offer to ICANN. And its services have been, to my mind, extremely expensive, not simply in terms of dollars but also in terms of the alienation that has been created between ICANN and the public.

Consequently, were I on the ICANN Board of Directors, I would work to replace Jones, Day, Reavis & Pogue.

#### **Reform of ICANN - Louis Touton must go**

Louis Touton left a partnership position at Jones, Day, Reavis & Pogue to become ICANN's Vice-President, Secretary, and General Counsel.

I believe that ICANN would be better off without him for the following reasons.

First, it has been my experience that it is a bad idea, both from a legal and business perspective to have the corporate counsel involved in non-legal operations. It tends to defocus both points of view.

Second, I have found that this person continues the tradition of condescension and uncooperativeness that has been the hallmark of ICANN's officers since the outset. ICANN needs people who build relationships without first looking at the social register or statements of net worth.

Third, in my personal interactions with this person, I have found him to be evasive and unwilling to provide even the most basic of information. I do not have confidence in him.

In addition, Mr. Touton has been an architect of ICANN's efforts to remove membership rights. See the section entitled "Full recognition of at-large members as "members" under California law." in this platform.

As a member of the board of directors I would have the right to rely upon the expertise of people such as ICANN's legal counsel - that is unless I felt that I was not receiving the degree and quality of advice that, in my opinion, I felt would be necessary for me to properly exercise my duties. Given my interactions with Louis Touton, I do not feel comfortable relying on his work.

#### **Reform of ICANN - Mike Roberts must go**

Mike Roberts is ICANN's President and CEO, and as such he receives an automatic seat on ICANN's board. He holds a "temporary" position - one that he has held since late 1998 - nearly two years.

There is nothing positive I can say about Mike Roberts. I have never had a positive interaction with him. I have rarely heard a statement from him that I believe is not laden with hidden agendas, unstated definitions, silent reservations, and secret conditions. And I have heard him proclaim, without any showing of evidence, that there was "consensus" for some policy decision when it was clear to me, and to others, that opinion in the internet community was far from agreement.

I have observed him at the helm of ICANN's business processes and believe that if ICANN had been a for-profit entity, it would have long since turned turtle and sank in a sea of red ink.

I believe that Mike Roberts has done an outstanding job of creating ill-will towards ICANN and distrust of ICANN's actions.

Under Robert's hand, ICANN's "staff" has assumed virtual control of major policy decisions. This staff operates nearly in total secrecy and without any form of public review.

ICANN doesn't need this kind of "temporary" executive.

In addition, Mr. Roberts has been a tireless architect of ICANN's efforts to remove membership rights. See the section entitled "Full recognition of at-large members as "members" under California law." in this platform.

#### **Domain Name Policy - Use existing laws, don't invent new ones**

ICANN, despite its denials, has enacted de facto worldwide laws that have as their practical effect the vast expansion of trade and service mark rights.

ICANN's rules supersede those of nations. That is wrong. ICANN should not be making law, much less should it be making supranational law.

(This lawmaking is especially egregious as it occurs without the participation of those who are most frequently the victims who lose rights and property as result.)

ICANN's UDRP is a case in point - most nations have laws that govern trade and service marks. And there are international treaties on the subject. ICANN's UDRP is an unnecessary addition to that body of law, particularly since it was created by a body, ICANN, that has no status as a legislature.

If a trademark owner believes that some domain name infringes on his/her rights in the trademark, there is an adequate body of law to determine if infringement has, in fact, occurred and, if it has, to accord relief.

Because it is duplicative of existing laws - laws that have been enacted by established legislatures - ICANN's UDRP should be eliminated entirely and absolutely.

Sure, this would cause trademark owners to sometimes have to travel to where an alleged infringer may be. And that is a burden given the worldwide visibility of domain names. However, our legal systems do evolve - the fact of evolution is built into their most basic foundations - and we ought to trust the slow, but sure, evolutionary processes of law to find a balanced result than to leap to an arbitrary law, such as the UDRP, created by a mere "private corporation" from processes that allow only one side, the trademark owners, to participate.

## Ted Byfield on the Election Process

BWG member on September 9: Okay, so let me see. 76,000 activated. On 15 August, we had 7407 activated, which was the last checkpoint, and the projection rate based on previous trend (icannnot.org) had 10,615.

Ted Byfield; <<http://members.icann.org/pubstats.html>> explicitly states that NA has 10,692 activated members, and that member-nominated NA candidates need 214 to qualify. I'd offer congratulations to those candidates who qualified, but i think I'll wait and see an official ack from ICANN that they \*did\*. But what I want to know is how many MALs weren't activated because of (a) 'technical problems' or (b) inadequate staff resources.

ICANN Announcement: NEW: Status of ICANN's At Large Membership: The nominations period is now complete. Over 76,000 applicants activated their memberships by yesterday's deadline. A new membership website

will be posted later today (California time), including information about nominees and a Question & Answer forum to facilitate dialogue. (9 September 2000)

**Byfield:** So if 158,000 applications justified crowing about how the MAL was 'an overwhelming success' (members-announce message, Thu, 10 Aug 2000 13:39:15 -0700), does the fact that less than 50% of them were activated render the process a 'miserable failure'? weinberg@mail.msen.com (Sat 09/09/00 at 10:25 PM -0400):

**BWG Member:** so if 158,000 applications justified crowing about how the MAL was 'an overwhelming success' (members-announce message, Thu, 10 Aug 2000 13:39:15 -0700), does the fact that less than 50% of them were activated render the process a 'miserable failure'?

**Byfield:** Not in ICANN's terms. They just figure that a majority of the voters were so happy with the NomCom candidates that they saw no need to participate in the member-nomination phase, and won't bother to activate until the actual vote.

According to ICANN: To vote in this year's At Large election, the deadline to activate your membership will be 8 September 2000 (midnight UTC/GMT). To activate, you must have your membership number, password (sent via email), and PIN number (sent via postal mail).

Byfield: I'm curious about this 'To vote in this year's election' bit. What's curious is that it does \*not\* say 'To be a member' (or however ICANN circumlocutes such things; remember that ICANN's web pages are all given a legal reading). What's the status of all these non-activated zombie applicants? is there any +/- to keeping them on the books in some way? Strains of Gogol's Dead Souls.....

**BWG member:** Did the \$144K [paid to Roberts] come out of the \$223 or is that in addition to? If the latter, that's a 9 mo. 'year' in which he and Darwin were paid \$367K. Can that be right?

**Byfield:** I'd guess the \$144K is a subset of the \$223K, based on this:

Date: Tue, 8 Aug 2000 10:16:49 -0700 To: tbyfield <tbyfield@panix.com> From: Mike Roberts <roberts@icann.org> Subject: Re: journalist query: 30 June 1999 Financial Statement Cc: Esther Dyson <edyson@edventure.com>, Louis Touton <touton@icann.org>, ajm@icann.org

This refers to the professional services agreement between ICANN and the Darwin Group which covers me in my capacity as startup CEO. During the first few months of the startup period, Darwin also handled some business expenses for ICANN before the corporation had its own accounting operation. These were reimbursed. Total payments for my travel and services are also shown on the payments to Directors report which is on the website in the same section.

The Darwin Group is a California corporation in which my wife and I are the principal stockhold-

ers. In addition to the disclosures here, this information is also contained in our submissions to the IRS and to the House Commerce oversight hearings last summer. - Mike

**Byfield:** does anyone have an archived copy of ICANN's loan-disclosure page?

It's changed: now, instead of a pseudo-disclosure listing one-year unsecured loans but no dates of execution, there's just a boiler-plate promissory note. And the year has mutated into a \*minimum\* repayment period:

<<http://www.icann.org/financials/tax/us/form1023-2att.htm>>

Response to Part II, Section 2

To date, ICANN has been supported by accreditation and application fees and by one-time, start-up contributions from corporations and individuals pending the establishment of a permanent funding mechanism based on fees paid by those using and receiving the benefits of ICANN's services as a consensus development body. A list of donations received to date is attached as Appendix 15. In addition, ICANN has obtained non-recourse loans from several entities for a minimum term of one year, to be repaid after a permanent funding mechanism is created. A list of the loans received is attached as Appendix 16. It is anticipated that ICANN's permanent funding will be generated from fees paid by registries and registrars (and possibly others) involved with the management and administration of the worldwide system of Internet names and numbers that use and benefit from ICANN's consensus development processes.

**Byfield:** other questions:

- if July 31 marked the end of the application period and Sept 8 marked the end of the activation period, when will another application period begin?

- if activated 'members' aren't really Members, what are zombies? is the \*only\* difference between them the privilege of voting in the MAL election? if so, that rather paltry distinction could have bearing on any legal proceeding bearing on (M)membership.

- are zombies still getting informational updates from ICANN? if not, when will they stop receiving them? And so on.

Date: Sun, 10 Sep 2000 16:13:37 -0400 From: tbyfield <tbyfield@panix.com> To: Mike Roberts <roberts@icann.org>, Louis Touton <touton@icann.org> Subject: query re loans' terms

I see that:

(1) and <<http://www.icann.org/minutes/minutes-26july99.htm>> says "a number of commercial entities have indicated a willingness to loan the Corporation funds for a one-year term at interest of 7% or below"; and

(2) <<http://www.icann.org/general/loaninfo.htm>>

says ICANN “executed one-year unsecured loan agreements” (but doesn’t give the execution date); and

(3) <<http://www.icann.org/financials/proposed-budget-04may00.htm>> speaks of “currently outstanding short term loans in the amount of \$1.025 million” and mentions possibly repaying them “ON SCHEDULE in August, 2000”; and

(4) <<http://www.icann.org/financials/tax/us/irs-letter-grant-28aug00.htm>>, dated 20 sept 1999 — i.e., later than (2) but earlier than (3) — say that ICANN has “obtained non-recourse loans from several entities for a MINIMUM term of one year, to be repaid AFTER A PERMANENT FUNDING MECHANISM IS CREATED.”

Why do the public disclosures about the loans’ terms differ from the disclosures about same made to the IRS in ICANN’s application for 501(c)(3) status? Thanks, as always, for your time.

**Ted Byfield:** If an organization’s application for 501(c)(3) status is approved, the org is required to disclose certain kinds of information—for example, its application. however, if its application is rejected, is it true that the org is not required to disclose this info?

The reason I ask is this:

(1) and <<http://www.icann.org/minutes/minutes-26july99.htm>> says “a number of commercial entities have indicated a willingness to loan the Corporation funds for a one-year term at interest of 7% or below”; and

(2) <<http://www.icann.org/general/loaninfo.htm>> says ICANN “executed one-year unsecured loan agreements” (but doesn’t give the execution date); and

(3) <<http://www.icann.org/financials/proposed-budget-04may00.htm>> speaks of “currently outstanding short term loans in the amount of \$1.025 million” and mentions possibly repaying them “on schedule in August, 2000”; and

(4) <<http://www.icann.org/financials/tax/us/irs-letter-grant-28aug00.htm>>, dated 20 sept 1999 — i.e., later than (2) but earlier than (3) — says that ICANN has “obtained non-recourse loans from several entities for a MINIMUM term of one year, to be repaid AFTER A PERMANENT FUNDING MECHANISM IS CREATED.”

In other words, if ICANN’s application for 501(c)(3) status had been rejected, the terms of the loan as described to the IRS wouldn’t have been required to be made public, right?

I think: ICANN is saying one thing to those who have no power to investigate its claims (the public), and another thing to the IRS, which \*can\* do so. can anyone give me a convincing explanation that’s more charitable?

Date: Mon, 11 Sep 2000 01:38:38 -0400 From: tbyfield <tbyfield@panix.com> To: bwg-n-friends@jetty.net Subject: Re: [bwg+] icann latest Mime-Version: 1.0 User-Agent: Mutt/1.2.5i

Sender: owner-bwg-n-friends@jetty.net Reply-To: bwg-n-friends@jetty.net

September 11, Byfield: after I zapped Markle for underfunding the MAL process, Laurie Baylin zapped me back:

Markle didn’t underfund the effort, but rather, gave ICANN double what they initially proposed — to provide that they include sufficient public outreach in their membership effort. When it was first called to our attention that there were system inadequacies due to unanticipated traffic to the registration site, we offered more in the way of matching funds — to encourage additional funder participation and express our continuing support of an open and participatory MAL.

The magnitude of ICANN’s misestimates in the case of the MAL is really remarkable. And the fact that ICANN was turning away grants from Markle in the very months when it was desperately negotiating loans from IBM, MCI, et al. makes ICANN’s true nature abundantly clear—and throws an awkward light on those loans.

If the technical systems underpinning the MAL were done properly — a \*big\* if—it isn’t really a question of ‘funding the whole thing,’ as you put it. The (trivial) infrastructure is in place, and the outreach has been done, so the costs would be expressed in staff time. And a good thing, too, in my opinion: as we have seen, the devil finds time for idle hands.

## Why I am Running for the ICANN Board by Larry Lessig

ICANN stands at a critical moment, and the decisions made by this body in the coming years will have a profound impact on the future of the Internet. At stake in this election is much more than who controls the directors of a non-profit organization, or a referendum on past actions of certain players in the process. No. This election — the first ever expression of self governance of the Internet — will determine its very constitution.

ICANN stands at a critical moment, and the decisions made by this body in the coming years will have a profound impact on the future of the Internet. At stake in this election is much more than who controls the directors of a non-profit organization, or a referendum on past actions of certain players in the process. No. This election — the first ever expression of self governance of the Internet — will determine its very constitution.

### THE CODE IS THE LAW

Over the past decade, I have studied and defended constitutions. Some of these constitutions were framed by legal texts - I am a law professor, and I teach and write about constitutional law. But the most profound of these constitutions is framed by code - by the software and hardware that makes up the architecture that constitutes the Internet. My aim over the past six years has been to understand the relationship between the archi-

ture of the Internet and its values, and to defend that architecture so as to defend those values.

Unlike most constitutions, the constitution of the Internet does not centralize power; it does not appoint central committees to decide policy for Internet users. Instead, built into its design is a commitment to decentralized power and to bottom up creativity rather than top down control.

ICANN is a potential threat to that architecture. As a body with the backing of the strongest governments in the world, but without the express sanction of an Internet community, there is a constant risk that ICANN will act beyond its limited mandate in ways that will compromise the initial values of the early Net. There is a danger that it will be captured by special interests, and then use its power to advance these interests.

### A LIMITED ROLE FOR ICANN

In my view, the challenge for ICANN is to keep its footprint small. Its role is to be a tiny, coordinating body for technical standards made elsewhere; its function is to assure those standards sustain and support the stability and diversity of the Internet.

Its job is not to become the trademark police; it is not to be the tool of intellectual capital; its job is not to set policy for the Internet generally (beyond defending the decentralized architecture of the original net); nor to create artificial scarcity, or choke points of power. ICANN must not be permitted to claim a mandate beyond the narrow tasks described in its charter.

ICANN must be accountable to that charter and the values it embodies. Its burden is to do this without becoming captured by powerful interests that would seek to manipulate ICANN’s position in service of their own limited ends.

If elected, I will work to ensure that ICANN maintains a vigilant focus on a limited, technical oversight and coordination role. I will staunchly defend free expression, privacy, open architecture, and security, by ensuring that ICANN stays out of making policy in these areas. And I will work to ensure that the original spirit of openness, diversity, and the free flow of Information — the original tenants of the constitution of cyberspace — are preserved.

I may not have the technical savvy of the Net’s original architects or of some who are candidates for this board. But I have studied and lived the process of designing structures to enable governance. That’s the task of a constitutionalist. And in my view, the structure of governance that we must preserve is the structure implicit in the original Net. ICANN may well do some good relative to this structure; but most importantly, it must do no harm.

Editor’s Note: Larry Lessig, author of Code and Other Laws of Cyberspace, is a nominee for the North American Member at Large (one of five being elected) of the ICANN Board. The above is the beginning of his platform found at <http://www.lessig.org>

## Executive Summary

### St Arnaud on OBG, pp. 1-9

We interview Bill St Arnaud, Director of Network Projects for Ca\*Net3 the world's first optical national R&D network. Arnaud describes the beginnings of a trend that may enable enough ISPs to peer directly with other ISPs so that instead of having to buy transit from core backbones for 80% of their traffic they may only need to buy transit for 20%.

The Optical Border Gateway Protocol (OBGP) is an experimental concept which, at this point, is unproven. The first enabling step is dark fiber and the availability of many dozens of wavelengths such that an ISP can then purchase its own wavelength and use it to connect to an exchange point several thousand miles away if need be. At such an exchange point you will be able to do standard peering. The advantage given you by OBG is that you can get to that exchange point without having to pay a transit costs to an upstream provider.

As the number of available wavelengths multiplies and the prices those wavelengths comes down, ISPs will find that they are able to buy wavelengths from their networks to dozens and perhaps even hundreds of exchange points some of which will be in other countries and on other continents halfway around the world. This is the first step. It's doable today and is beginning to happen as Williams and others are now selling wavelengths to ISPs.

With OBG any optical switch in an optical network can be treated as an internet exchange point such that autonomous ISPs can interconnect and peer with each other anywhere along the network. Now this has some profound consequences. If you let users at the edge control the routing and topology of the network through control of the ports on the switch that means the carrier in the middle will be very limited in how it can optimize and manage the wavelength routing in the network.

BGP has an options field and a number of proprietary products have been using that field for their own special purposes. They are proposing to use the options field in BGP to turn networking upside down. Today when you connect an ISP to an upstream ISP, the first thing they do is to install the physical fiber. Then they put in the link layer which can be ATM or SONET. Next they establish IP connectivity and finally BGP connectivity. They are saying let us reverse that whole process.

So first they would establish BGP peering. They would say I want to peer with Gordon Cook. They have your AS number and can start a BGP session where they instruct their router to connect us. It is the router then that establishes the physical connection between us. Right here in Ottawa most of their institutions are going to put in either for four or eight wavelengths. They can set up their wavelengths set up to go to whichever universities they choose. Then if someone says

let's set up a BGP session with Gordon Cook, a router will take one of those wavelengths and steer it towards a connection that will peer with your network.

In their Quebec university-owned dark fiber network they will do a proof of concept that will demonstrate these capabilities over the course of the next couple of months. They will put together and test a very simple version of the protocol. They already have a number of industry partners involved in the project who have indicated to them that they're ready to take the concept commercially and go with them into the IETF.

For one time investment as little \$10,000 it is possible to get wired with dark fiber that will last for 20 years. The biggest single expense for institutions controlling their own dark fiber is internet transit cost. If OBG becomes usable the amount of transit needed for purchase should decrease to between a fourth and a fifth of what is necessary to buy today. What works to reduce transit expense for universities will also reduce them for small ISPs that could use wave lengths and OBG to peer with each other.

Arnaud expects that in the future there will be three parallel networks in existence. In fact he would suggest that policy not to try to converge all telecommunications into IP because doing so will increase the cost of IP. The residential telephone market will continue to provide delivery of telephone by twisted pair. There is nothing wrong with a voice-over copper which in fact works very well. A second network may well be broadcast video over coax which again works very well. There is no really compelling reason to deliver broadcast video over IP. I believe there will be a third network dedicated to only IP and this network can under some circumstances carry voice or video or both.

When asked what must be done to continue the scaling of backbones, Arnaud comments: "There are probably three possibilities. I am not a fan of the MPLS approach. MPLS was a technique designed to cope with shortage of bandwidth. Now with the ability to buy wavelengths of light on fiber bandwidth is not an issue. Consequently micro engineering the network is probably not necessary today. Also it does not look like we can build routers big enough to aggregate all the traffic. We will require optical "cut thru" or "bypass" circuits. The challenge that faces is who controls the cut through circuits - the carrier in the middle or the customer at the edge."

"Now, doing it this way is not building a traditional circuit switched model. Circuit switched models imply that for every flow you must start a new circuit. The service will set up a switched circuit and send a web page to you, and then switch to another circuit and send a web page to someone else. This will not scale. Having direct peering with each other with wavelengths and doing bypass around other ones may achieve the same ends. In the standards bodies, you have three approaches. One is called the overlay and this is basically the circuit switched model. This is being promoted by under the label of ODSI. Another one is called peer networking. This is where all the wavelengths are treated like MPLS tunnels. The third approach is ours. In this we

say let BGP be the controlling mechanism and let decisions be made more upon the lines of traditional Internet direct peering relationships."

### World Wide Packets, pp. 10-14, 26

We interview Bernard Daines the CEO of World Wide Packets. Daines explains that the explosion of dark fiber over the past five years has been fueled by fiber networks laid by cable TV companies and by utilities in addition to the more familiar carrier networks. In many place in the US utilities are bringing fiber to existing homes and especially to developments of new homes. The dynamics of DWDM and Gigabit Ethernet mean that the cost of bringing fiber to new homes and of bringing telephone, video services, and internet as well has fallen low enough to enable a homeowner to pay for the services that only a few years ago a wealthy corporation could afford.

World Wide Packets is re-engineering gigabit Ethernet switching equipment for use in fiber to the home environments. It is driving down the cost of such equipment to the point where a hub and spoke distribution system for 100 homes in the form of a central community distribution system and subscriber distribution units for each home would cost a total of about \$1200 per home. Their equipment could be used to provide telecommunications services to customers of companies like VDN in Montreal about which we wrote in our last issue.

In addition to the Subscriber and Community Distribution Units World Wide packets provides its customer with a Network Management Unit (NMU) that is basically, a network management software application. Fiber owners or network administrators would use the NMU to provision services to their subscribers, manage network operations, and define service levels. It provides a GUI interface, which can be Java-based or Web-based, to control the CDU and the SDU—and to give the cable or utility customers of World Wide Packets complete end-to-end management of their entire network. In short World Wide Packets is enabling the creation of micro-telecommunication companies that by finding community market niches can use the cost efficiencies of fiber and gigabit Ethernet to compete against the older and far larger and more wealthy but technologically backward ILECs.

### IPv6 Deployment, pp. 15 - 18

We summarize arguments on IPv6 on IETF and Nanog list from late August and early September. It is pointed out that advocates have alienated ISPs whose help they need.

### ICANN Elections, pp. 19-24

We publish documents relevant to the campaigns of Karl Auerbach and Larry Lessig as well as evidence of Ted Byfield's efforts to find out how the code for the membership campaign and elections servers was acquired.

Continued from page 14

<http://www4.nationalacademies.org/cpsma/cstbnsf/44bf87db309563a0852566f2006d63bb/d4a118651b44c3128525693e0053b588> Open Document

This URL will bring you to the White Papers Prepared for Broadband Last-Mile Technology Project page, which contains links to Acrobat versions of the following papers:

George Abe, Palomar Ventures "Factors Influencing Investment in Residential Broadband Equipment and Services - a Venture Capital Perspective"

Ian F. Akyildiz, Georgia Institute of Technology "Broadband Satellite Networks for Last Mile Technology"

Andreas Bechtolsheim and David Cheriton, Cisco Systems "Ethernet Broadband Networking"

Edward Ciesla, Flack & Kurtz Consulting Engineers "Broadband Services to Rural Western Massachusetts"

Tom Dennett, Harmonic Data Systems "Last Mile Connectivity Utilizing Fiber Satellite Solutions"

Sylvie Doucet, Planned Approach, Inc.

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

"High Bandwidth, Applications, and Economic Development: Lets Tie It Together!"  
<http://www.plannedapproach.com/community.htm>

Richard Esposto, Western Integrated Networks "Getting Tele/Tech on Local Government Radar"

Einar Edvardsen, Telenor R&D "Broadband Access Over Inverse Multiplexed Copper"

Aura Ganz, University of Massachusetts "Technology Developments for Quality Multimedia Delivery for Residences: Coupling of the Broadband and Home Network Technologies"

Henry Kilpatrick and Paul Baker, Georgia Institute of Technology "Regulatory Issues, Pricing, and Access to Public Utility Right-of-Way"

Jose-Marie Montpetit & R. Deininger "The Use of Satellite Technology for Last Mile Broadband Access"

Jose A. Pozas, Telefonica I+D "Deployment of Multimedia Services to Residential Customers"

Amnon Ptashek, EDSL Networks Inc. "Residential Internet Ready Buildings (IRBs)"

**Battle for Cyberspace-  
How Technology and  
Political issues May  
Affect your Internet  
Venture cost of \$695  
and now available.**

### **Subscription Rates**

Choice of either ascii or Adobe Acrobat (PDF) format 1. Individual; College or University Department; or Library; or Small Corporation - \$250 2. Corporate - (revenues \$10 to 200 million a year) - \$350 3. Large Corporate- Revenues of \$200 million to \$2 billion per year - \$450 4. Very Large Corporate- Revenues of more than \$2 billion per year - \$550  
Site License: The right to distribute ascii and PDF via email to all employees of corporation. 5. Small corporate: \$450 6. Corporate: \$650 7. Large Corporate: \$900 8. Very Large Corporate: \$1150 . Site License Distribution via intranet web site \$400 a year additional. See [www.cookreport.com](http://www.cookreport.com) for more detail

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