



The COOK Report on Internet



© Cook Network Consultants

Bandwidth Commoditization Technology Enablers Sycamore Networks Focuses on Intelligent Optical Networks Using Software to Control Light Paths Web Based Interfaces to Give Users Ability to Set up and Tear Down Circuits Permitting Bandwidth Management

Editor's Note: Jeff Kiel, has been Vice President, Product Marketing, Sycamore Networks since the summer of 1998. His responsibilities include product management, strategy and marketing of Sycamore's technology. Prior to Sycamore, he was at Ascend Communications and was responsible for their ATM frame relay switching products. Before Ascend, he had the same responsibility at Cascade Communications. Prior to joining Cascade in 1996 he was with Bellcore for eight years and BellSouth for three years. We interviewed him on April 7th and followed up with more question and answers in mid June.

COOK Report: In order for optical networking to reach its potential will it have to develop technology that will result in a commodities market in bandwidth as a by product? Also if standards are involved, are there some issues with people not being comfortable with a "vendor-driven" solution rather than a vendor independent?

Kiel: Sure. Well, there's a couple of questions in there. Let me start talking about the first one, which is, given the realities of where carriers need to take their networks and with the fact that a bandwidth commodity market is coming, you are certainly going to wind up putting a new set of requirements on the underlying optical network.

Defining the Optical Networking Market

Sycamore develops hardware and software that carriers use to light the fiber and create high speed end-to-end services. Carriers have optical networks today. They've been building optical networks for several years. Those optical networks have been comprised

of the fiber and then of DWDM that multiplies the bandwidth available from the fiber. Then they built out SONET or SDH rings on top of that. The SONET supported a set of technologies and products that evolved when the applications were low-speed applications, predominantly voice. Also the infrastructural services supported by SONET had been very long-term — 15, 20, and even 30 year services.

But now, with the revolution driven by the Internet and DWDM, the applications that drive this network are changing. Instead of having lower speed, very long-term voice dedicated services, you're now having high speed, relatively short duration data services. And bandwidth trading, as part of a bandwidth commodity market, is just going to exacerbate the transition to the kind of marketplace where connections need to be set up and torn down, really in just a matter of seconds, if not milliseconds. Such a network needs to be able to run in the most efficient manner possible.

COOK Report: You started off by emphasizing the carriers. But is it not a company that is primarily an owner of dark fiber, for example a company like your best customer, Williams Communications, which probably has the most significant interest in the commodities market at this point? Would that be a reasonable statement?

Kiel: It is both the carriers and the service providers that light dark fiber and sell services. They're the ones that are being impacted by the change in demand and the change in the service applications.

COOK Report: But aren't we talking about carriers in the sense of fiber owners, not that AT&T doesn't own fiber. I think they own

Volume IX, No. 5, August, 2000

ISSN 1071 - 6327

quite a bit, but certainly they haven't trenched the amount of fiber that a Level 3 and a Qwest and a Williams has recently.

Kiel: Yes. You're absolutely right in that, what the impact is going to be on those fiber owners and the differentiation between carriers is how do they light that fiber? What is that technology that they bring to bear to light that fiber?

COOK Report: Right, and the SONET technology is on the, what's the right word? It's going to be less and less used.

Kiel: Absolutely, because that SONET technology was originally designed and implemented to deal with low-speed services, voice services, long holding time, long duration contract.

And so with users not being able to predict

On the Inside:

| | |
|-----------------------|--------------|
| Sycamore & Bandwidth | pp 1- 6 |
| AIG Telec & Bandwidth | pp. 7 - 12 |
| In-addr.arpa | pp. 12,20,24 |
| Airwire.net | pp. 13 -15 |
| ICANN Elections | pp. 16 - 18 |
| Silicon Valley | pp. 19 - 20 |
| Juniper vs. Cisco | pp. 21 - 22 |
| Executive Summary | pp. 23 |

when and where they're going to need traffic and high bandwidth, as well as the holding times are going to be reduced from years down to months, down to days, literally down to hours, the existing networks that are in place are not geared to respond to that. It's not that the existing networks are bad, it's just they were never designed for that. These applications were never envisioned when all the vendors had started to develop their SONET equipment, their SDH equipment, as well as their DWDM equipment. That equipment was developed for a very static, predictable growth environment.

COOK Report: So you guys are providing equipment that is going to be used in the changed situation. Would it be appropriate to begin describing what it is now that you bring to this table?

Kiel: Yes, and that's exactly what we're focused on: this changing environment of the applications and the user needs and bringing a next generation set of equipment to light that fiber. We refer to this as Intelligent Optical Networking, something that comprises a new class of optical equipment. We started to see some of the initial introductions of that equipment in 1999. Carriers started to begin implementation in 1999. And it's going to be an evolutionary process of them moving from their existing SONET ring-based networks to more of an optical, mesh-based network.

The key differentiator we see is the software control of the light paths. If you take a step back and you look at the optical network, you realize that you have to create the light. That's your lasers and your amplifiers and your filters. You've got to be able to package that light, meaning you have to be able to hand it off to a customer in a consumable format, since people do not consume raw wavelengths. You also need to switch that light, you need to get that light from customer A through that network to customer B. And then you've got to be able to manage that light.

And so the last two are going to be the key differentiators for these new set of products that come out, because ultimately, we all go to the same wells for the lasers and the amplifiers and the filters, because that has now become almost a commodity market where there's companies that specialize in those. And so companies like Sycamore can rely on companies specializing in lasers and filters and amplifiers and they will continue to invest and push the envelope, whereas what we can do is focus on how do we assemble our system, what's the architecture? But, more importantly, what is the software that we bring to control that architecture?

But if you think about the optics, every day lasers get cheaper, they get stronger, they go

greater distances and they run faster. A carrier is going to be introducing new optical technology as quickly as it makes economic sense into their network. And the only constant that exists throughout their network as they move to adopt this rapidly changing optical technology is going to be their network management software.

COOK Report: Well, how would you describe what you are bringing to the table? What needs to be done that would give the software management that would allow these light paths to be traded in a commodities sense?

Kiel: Sure. To answer your question in a more general sense, the other vendors are suppliers that had come into being during the days of DWDM. They focus on fiber multiplication and on the physics aspect of DWDM. Namely, how do I put a few more waves on a fiber strand? Instead of creating 50 virtual fibers, they want to create 100 virtual fibers or 150 virtual fibers.

COOK Report: And those virtual fibers are just essentially sold as a wave from point A to point B. It's just there and the problem is that you can't quite massage to the extent that your philosophy says you'd like to be able to? Is that getting warm?

Kiel: Yes. That is correct, and what they're doing is focusing on taking one fiber, creating multiple fibers. They are looking at these networks really as segments.

COOK Report: So a segment in the sense of a physical link from point A to point B, rather than something that is more dynamic and changing?

Kiel: Exactly. Because what's happening is that the role of the optical network is changing. The optical network today, fundamentally, its role is one of plumbing. It provides pipes. It's something that is hidden from the user. The customer doesn't see it, they don't know that it exists. It is a set of technologies that is built as plumbing, or as pipes.

What's happening is that the users have changed what they want to buy. They want higher speed services and they now want to buy them in hour increments instead of 30 year increments. What you've got to do is take a view of the optical network as a series of paths, actually as a series of end-to-end paths which in turn can be seen as a set of concatenated pipes, or set of concatenated links. It's the management of a network on an end-to-end basis, from an ingress point to an egress point, which is the biggest challenge. So what you see Sycamore focused on is viewing a light path through this very complex network in exactly the same way that the data network views a virtual circuit.

Kiel: In a network you have to be able to very quickly establish a high-speed service end-to-end. In order to do this you select an ingress point, select an egress point, a set of service attributes for that connection and then let the network set up what you have chosen. That is the process by which you activate service. And you see a lot of folks now focusing on the speed of provisioning, because that is really the level of competitive service that carriers need to provide, particularly for this bandwidth trading market.

COOK Report: How might that work in a hypothetical situation or two?

Pushing Network Intelligence Downward into Each Network Node

Kiel: What we actually have done is to push the intelligence of the network down as low as it can go, actually into the network nodes themselves. And that is a completely different philosophy than most vendors that focus on centralized intelligence. You will find, in the business model of these folk, a centralized point and a centralized database that has awareness of the network and will determine what's the best way to route a connection from ingress to egress. What we've done is distributed that intelligence into all network elements. So the network elements themselves are self-aware of the network.

COOK Report: And how is this accomplished? I mean, physically, what are some of these network elements? Are they certain kinds of Sycamore physical devices with Sycamore software?

Kiel: Yes, these include the Sycamore switches, the 16000, for example in which we've implemented IP-based routing.

COOK Report: The "16000" is the name of the switch?

Kiel: Yes. We've implemented an IP-based routing and control mechanism as a means to control wavelengths inside the network. And this is where OSPF, which is what John Moy works on, along with MPLS, is used inside the optical network to control and manage the wavelengths end-to-end.

COOK Report: How does this switch then differ from a device that a Ciena might have to sell? Is the fact that you've got the software down at the node level inside the switch, does this mean that there's some CPU and memory associated with each physical switch that wouldn't be in another vendor's solution?

Kiel: Well, what you have is a number of vendors out there that are working on developing these switches. A lot of folks tend to get focused on how large the switch is from just a size perspective as well as how many ports does it have. Whereas what we've seen is that the real differentiation is going to be on how large can these networks actually scale.

Let me give you kind of a corollary. If I go back to my data communication days at Cascade, there was a time when the different ATM switch vendors were looking for bragging rights for who had the first OC12 interface on their ATM switch. It really wasn't all that long ago.

As it turned out, back then the time difference between availability of an OC12 interface was going to be measured in weeks, and if not in weeks, then a matter of months. However, what ultimately determined which communications vendor was successful or not was not whether they had OC12 on their switch or not. It was how many switches could they support in a network? Could they scale the overall size of the network into the hundreds of switches?

COOK Report: And the answer to this would depend on what kind of network management software the switch vendor had. Right?

Kiel: The software within the data comm network became the gating factor to everyone's ability to scale the data network. And that ultimately determined which vendors were successful and which vendors failed.

We see the same exact thing happening here. This network is going to grow bigger, and it's going to grow faster than anything any of us have ever seen before.

So while people are still getting hung up on how many ports does your switch have, how many gigabytes? What's ultimately going to matter is how many switches, how many thousands of switches, how many millions of connections, how many tens of millions of users, can be supported on this network? Once again, networks' limit in size and profitability is going to be determined by its software rather than its hardware.

COOK Report: Would you agree that there seems to be an analog here to the Isenberg's ideas about this stupid network and moving the intelligence to switches that are more at the periphery?

Kiel: The only way to scale is to distribute the intelligence down to the network elements themselves. If you have one central brain in your network, you face several issues. One is, certainly, if anything ever goes wrong. And the second thing is that central-

ized brain can never scale fast enough so that it can deal appropriately with everything that's going on inside the network.

COOK Report: So you must have to have some sort of console or device. Something that takes the input from the hundreds or thousands of switches that you have, and it's almost like a master screen that aircraft controllers might look at that allows them to see the entire network at one time. In the human interface to the network management software does there have to be some kind of a centralized point? Yes or no?

Kiel: Yes. And one of the important things we've done with our software is automate a lot of the manual processes that exist in today's optical networks. In these networks, there are a number of offline databases that need to be maintained. You need to manually configure and preplan networks, as well as have people out in the field actually set up these devices and provision them.

But what we've done is put into software and automate that management process to make it very, very simple to bring these networks up, as well as make it very simple to provision services. To give you some specific examples, when you take one of our switches and put it in a network, the first thing it does is figure out who its neighbors are. And it will understand and learn, on its own, what the overall network looks like. It will report upward to a network management system, which provides an overall map of the network. So the network itself is the database, where the network always holds the most accurate information about its state. This data presents itself up to a centralized GUI for the network operator to look at.

Now, to establish connections from the network operator's perspective, it's very, very simple. They pick one port on one switch, they pick another port on another switch and then they give a set of service attributes, such as what speed is my service? What type of protection do I want on that service?

COOK Report: In other words, if I'm at this screen, I may want to do something with a switch in Seattle and another one in Herndon or Vienna or Virginia, Washington, D.C. suburb. So the first thing I could do would be to select the geography involved?

Kiel: Well, what you do is, it's as simple as point, click and done.

COOK Report: And then when I've selected the geography, it presents me with a screen of attributes that I want to establish for the circuit that I want to create.

Kiel: Correct.

COOK Report: And, these attributes would

be what kinds of things?

Kiel: Some examples include, what is the size of the circuit? What's the desired protection for the service I am establishing? Do I want assurance of millisecond restoration? Or can I do without protection? Can I permit my service to be preemptable?

COOK Report: Because if I say I want protection, that protection presumably is protection against a physical break. Now, in order to get that, I'm going to be committing some resources elsewhere on my network, yes?

Kiel: You may or may not have to commit resources. The answer will be determined by the speed of recovery. To get the fastest speed of recovery, you must dedicate resources on another path through your network for the user in question. Whereas if you start to share, an alternative is that you set aside a fraction of the available capacity. This becomes your protection bandwidth. And what's interesting is that your protection bandwidth is not fixed in place, it's actually fluid and floating throughout the network.

Readjustment of Network Resources

I equate this situation to having a puddle, that's the amount of capacity you have in your network. What you want to do is have a smaller number of drops in the "puddle" set aside in the event that something, there's a fiber cut or there's some other type of network fault condition. Imagine that surplus protective capacity as drop of oil placed on top of your "puddle."

That "oil" is just going to float around, like a sheen, on top of the water. It's not going to stay in the same place. Success will come to the carrier and ultimately the vendor that can most efficiently make that drop of oil as small as possible. Those are going to be the ones that are successful. If I am competing against you, I don't want to have to build a network twice the size of yours to support the number of users.

Ultimately, you want to have a network that is in the lowest cost service position. But you also want to have a very rich software base that allows you to bundle in other services on top of basic connections. Such services are there to allow the service providers to continue to enhance their service offerings and maintain their margins. So one example of future needs will be not only how quickly can you provision, but also deprovision a service.

COOK Report: So, you're saying a small

amount is kept in reserve and if you've got a problem, your network management software can say, okay, there is this is the kind of problem from this point to this point. It can send some signals out to commandeer some of the surplus bandwidth and start routing it in a different fashion from where it had been going before to bring it momentarily into the breach that might have been caused by whatever condition had happened.

Kiel: Right and what happens is the switches do that themselves automatically. Because if you can imagine having to go back to a centralized location and have that centralized location sort it out, it's not a scalable solution. It's going to bog down.

COOK Report: So when you're establishing the parameters of the switches at various places in your network, do you have a choice where you say what is the surplus ceiling that this region of my network may have? In other words, if I've got a problem here, do I wish to allocate to these switches the ability to immediately grab a certain amount of bandwidth or does it work a little bit differently?

Kiel: Well, what happens is the carrier has to decide how big they want that drop of oil. So we can work with that to make some recommendations and then that's essentially referred to as traffic engineering, which is they just have to plan, they built a network of capacity X and then how much of X will they sell? If it's a SONET network, they can sell up to X divided by 2.

COOK Report: A lot less than with a non-SONET.

Kiel: And then with an optical, mesh-based network, what they can do is get very, very close to full network utilization. And then what the carriers then have to decide is how close do they go to 100% utilization? Because that is dependent upon the mix of their services. So what you can now do is start mixing protected services with unprotected services with preemptable services. And that's a stark contrast to SONET, which every service is protected, whether you want it or not.

COOK Report: Okay, and this is a whole new approach to a kind of Quality of Service. You can perhaps use this bandwidth variability and network management capability to achieve QOS as a part of an overall network management system, where you have an environment that is far more flexible, and more quickly flexible, than your standard SONET-based environments of today.

Kiel: That's exactly correct, but there's one additional thing worth noting here. While these concepts we're discussing, are abso-

lutely brand-new and revolutionary to the optical portion of the network, they are old-hat to the datacomm part of the network.

COOK Report: They're new and revolutionary to the optical portion of the network? Why? Because it's always been based on the ideas of much more static, SONET types of provisioning?

Kiel: That's exactly correct. When you think of a datacomm network, you think of a very dynamic environment of connections coming up, connections going down. The notions of Quality of Service, service level agreements, and the ability to partition the network based on virtual private networks were all well accepted and implemented ideas that were born out of building the public data network.

Now, these same exact principles are being applied for the optical network. So on the one hand, they are revolutionary in the sense they've never been used in these lower parts of the network before.

COOK Report: Are they going from the TCP/IP level down to the network level, then?

Kiel: That's exactly right.

Role of Optical Domain Service Interconnect

COOK Report: How do the ODSI developments relate to all of this?

Kiel: Optical Domain Service Interconnect or ODSI is an initiative that Sycamore and a large number of other vendors started back in January of this year. What the ODSI is focused on is the interface between the optical network and the IP network. For the last few minutes we've been talking about some of these new technologies that are being applied to the optical network, where they can be used to set up and tear down connections on demand.

At the same time, within the IP network you have very sophisticated and advanced traffic engineering capabilities, where they can now sense when they need a new connection. So what we're doing is marrying those two technologies to allow the IP layer to essentially call up optical layer bandwidth on demand. It's almost like an optical dial tone.

What happens is that we've provided the ability for a human user to point and click and set up circuits and tear them down, but once those mechanisms are in place, what you really want to do is get the humans out of the picture. Because it's the routers, it's the other types of data devices, they're the ones that understand the traffic patterns and

the traffic flows. Let them now set up and tear down connections when and where they need them.

ODSI is comprised of three groups; a group of service providers, equipment providers that have equipment in the IP domain, and a set of equipment providers that are building equipment within the optical domain. We are looking to establish a way for the IP layer to set up optical light paths on demand. We've just completed our second meeting. In fact that was yesterday, (April 6) I was out in Chicago for that. Our goal is to be able to demonstrate multi-vendor inter-operability of this concept by the end of this year.

COOK Report: Very interesting. In the earlier part of the conversation, I had an idea that with your overall network management and a GUI interface, I could provision an OC12, say, between Seattle and Washington, D.C. and tell it to start at 6 o'clock this evening and run till 6 a.m. the next morning. And it's just there, at the physical network level, presumably. Or at least at level 2, at the network level, rather than at level 1, which would be the physical. Now, if these assumptions are correct, what do these developments that you just outlined bring to the capability of someone who might want that physical path a day from now for a period of twelve hours?

Kiel: Let me answer your question by explaining how ODSI is envisioned to work. Based upon traffic flows and upon IP Quality of Service within the IP service layer, they may decide that they need to establish some additional circuits between a router, let's say, in Chicago and one in New York. So rather than having to have one human pick up a phone or fill out a service order, send it over to the carrier, have the carrier implement it and have the carrier's operator point and click the circuit, just let the router signal directly to the optical network. It will sense that, jeez, I need a little bit more capacity between these two city points. I'm going to set that up right now.

COOK Report: So if you're taking a fairly big pipe like an OC12 and you're using it for a given period of time from point A to point B, are you saying that purpose ODSI is to give you some increased intelligence? And, as a consequence, while the circuit is open, you can slice and dice and get much more efficient use out of a fairly large stream of bits than you would be able to do without this ability?

Kiel: Right, because let's say those two routers have an OC12 and then for a momentary point in time, let's say it's during a peak business time of data backup, what they want to do is, instead of running an OC12, they want to have an OC48 between themselves. And they can just do so, by signaling the

networks. Saying to the optical equipment beneath them set me up a connection from here to there, because I want an OC48 and I want it now.

COOK Report: Give me between these two points a little bit of that “oil sheen” for a minute or two that’s floating on that metaphorical puddle?

Kiel: That’s it. And then when it’s done, it basically tells the layer 2 optical network, I’m done, and it hangs up. Much like someone makes a phone call. And so that’s what the ODSI is focused on and what we’re doing is using IP-based techniques designed in such a way as to allow the routers to control the underlying optical network. What’s nice about this is the routers themselves are already using MPLS as a means of controlling IP traffic within their own networks. We intend to use those same control techniques to allow the router to talk to the optical network to set up and tear down connections whenever they need them.

COOK Report: Fascinating. Who came up with this idea? Was it primarily Sycamore who led in this development? Can you put these ideas in the context of last year or so and more sharply focused and help me understand from a larger industry perspective, rather than Sycamore, per se, how these issues might be used. I understand the benefits, but if I’m jumping outside of Sycamore momentarily and into Ciena or to some other company, what might be some of the anxieties?

Kiel: Well, I don’t think there’s any anxiety there. We started seeing tremendous interest across the industry. You mentioned Ciena, Lucent, Nortel, all those companies are participating in this effort.

COOK Report: Is there a mail list involved in this or not?

Kiel: There’s a web site and on that are the member companies and you’ll recognize a large number of those, as well as a lot of young odsi companies. The url is www.odsi@coalition.com.

COOK Report: Where does the OIF come in?

Kiel: The OIF is the Optical Interworking Forum.

COOK Report: Oh, okay, that’s the one that’s been in effect for some period of time, right?

Kiel: That was established a couple of years ago. The ODSI is basically a coalition of interested service providers and interested vendors looking to push this idea and get it defined and get it implemented. And then, what we’ve been doing is, after every meet-

ing, take all our results and take them right to the OIF. So we’ve been working very closely with them. A lot of the same people that go to the OIF come to the ODSI, we’ve got a large number of people contributing to developing the specifications in both forums. The role of the ODSI is to prime the pump, bring focus to this industry problem and then get an initial set of specifications working.

Commodity Bandwidth

COOK Report: To bring this discussion back to the subject of commodity bandwidth, if I understand what you’re telling me and you’ve got fully implemented ODSI capability, this would mean that I could get more efficient use out of bandwidth that I was purchasing as a commodity and therefore that I might be willing to pay a bit more for it.

Kiel: Yes.

COOK Report: Then, can I ask you to summarize for me what you think is needed for commoditization? Obviously the benchmark and the uniform industry contract are needed. Are there any of the issues we have been discussing that impact on what is needed to get those established? What does have to happen between now and that hypothetical point in time a year, or 18, or 24 months from now, when we may have an automated commodities market. One where we could log in with a GUI interface and see what the available price on an OC48 from city A to city B for some specified period of time is?

Kiel: I think what will ultimately speed up or slow down the creation of the commodity market is the following: if you think about what has to happen, at the end of the day, somebody, somewhere has to have a connection established and then a connection torn down. It has to be done very, very quickly. It has to be done at a very precise moment in time.

And so what has to happen is that the underlying network needs to be re-architected and rebuilt to give that type of control to the carrier. And what the carrier will do is take that very nimble, very agile and very flexible network and use it as a tool to allow them then to move into the bandwidth commoditization and the bandwidth trading market. Therefore the biggest challenge within the bandwidth trading, bandwidth commodity market is going to be having a network that can respond to these new requirements.

COOK Report: And in order to have a network that can respond, the ODSI definition of, standard criteria has to be further refined and accepted by most of the players in the

field, such that their devices are performing in accordance with this? And if presumably most of the people have the network capability and the devices that can do this, then you have an overall environment of the first time where it is presumably fairly easy to overlay that kind of automated, highly flexible trading system on top of it?

Kiel: I would say there’s two pieces to it. The first is where it’s actually still going to be the service provider that’s going to be setting and tearing down connections in response to the bandwidth trading and that does not require ODSI. But ultimately what’s going to happen is, the actual end users themselves are going to be setting up and tearing down connections by themselves. And that is what is going to require ODSI.

COOK Report: Okay, that opens one maybe last question in that I was talking to Eric Rabb last week of AIG Telecom, who had been to that late march Washington DC bandwidth brokering meeting. He was telling me his perception of some MCI people that he had seen there. He thought that they were perhaps less than — and I don’t mean to pin it on any particular vendor — happy about the situation in which perhaps their own customers could go into this commodity market and buy their own DS3, perhaps from point A to point B?

And it sounds like that’s what you’re getting at. If I understood what you just said, at the carrier level we have the ability, once the benchmark and contract is established, to buy a rather large bit pipe from a point in time A to B and set it up and tear it down. However, you are also talking about a set of technologies that will move the ability to do this from a small number of rather large companies that purchase large amounts of bandwidth down the food chain to enable the customers of these large companies to buy their own bandwidth. If these companies could acquire their own bandwidth pretty much in real time as they needed to do so, the outcome would be pretty profound. Right?

Kiel: Ultimately, that is where these networks are going. Because of the realities of data traffic, you know, when someone clicks on a website and it has some high quality streaming video, they’re going to need a connection right then and there at that point and time. There was no advance planning. You can’t predict when and where you’re going to need bandwidth. So these underlying networks are going to be need to be able to react to what’s now becoming these new applications of I want bandwidth now, I want very it high quality and I only want it for five minutes.

This is what’s causing the underlying optical network to be completely rebuilt. And if

you look at the optical network itself, it really is, it is the last vestige of Old World technology and it's the only part of the network that has not been rebuilt and re-architected given the realities of the data communications traffic. So we're just at the cusp of what is going to become a very exciting time.

COOK Report: Absolutely. So to finally conclude, then, if you look ahead at the next 12 to 36 months, and from the point of view of commoditization of bandwidth, what do you see as the major accomplishments that have to be achieved over some period of time from right now into the future?

Kiel: One of the major benchmarks to look at is how quickly can you establish the service from the time you decide you want it? Also how much time do you need to buy the service you want.

COOK Report: So would it be correct to say that once you have an industry benchmark and a uniform contract you can begin to trade in a non automated way. By voice telephone, one could then begin to set up a contract for an OC48 from September 1st through September 30th of the year 2000 between player A and player B, sign on the dotted line, fine, done. What the next two years will bring is the ability to do that sort of thing much more quickly, in much smaller pieces and in a much more automated way?

Kiel: That's exactly right. The metrics to measure against these underlying networks will be how quickly are they going to be able to respond and then what is that smallest increment of time in which a service can be sold?

COOK Report: Okay, terrific. Are there any issues that you see at this point and any contributions that Sycamore may have to make to the resolution of those issues involved in the definition and establishment of a uniform industry benchmark and contract?

Kiel: Well, I think as far as establishing the benchmark and contract, that is something the marketplace will need to sort out and the carriers will be focused on doing that.

COOK Report: I mean, there's no major technology hurdles or issues there?

Kiel: Well, what I'll say is that the concept is well understood. We're in the mode now is one of just pure execution. Because while these concepts, again, appear to be revolutionary, they're only revolutionary to this part of the network. These techniques of being able to rapidly set up, tear down connections have been implemented within the frame relay and ATM networks. So here is what we see happening now and what our contribution will be: bringing those types of

management and control approaches to the optical network itself and giving carriers the ability to build a network that can respond in real time to their needs.

Postscript (June 2000) - Standards Issues

COOK Report: How does the following OIF report below from your web site, relate to the ODSI work? Has there been a third ODSI meeting? If so what happened?

OPTICAL INTERNETWORKING FORUM A new Signaling Working Group has been formed to focus on signaling between the electrically and optically switched domains. An interoperable User to Network Interface (UNI) is expected during 2000. The OIF's Architecture Working Group has completed a draft Architecture Requirements document outlining future work. The OIF's Physical & Link Layer Working Group completed key internal system interface specifications for OC-48 and OC-192 data rates. The SER/DES Framing Electrical Specification for OC-192 formulates a common interface between framers and serial/deserializer parts used to implement OC-192 and STM-64 interfaces. The SPI-3 Packet Interface for Physical and Link Layers for OC-48 provides system designers with a common Packet Over SONET (POS) Physical Layer interface, providing a versatile bus interface for exchanging packets within a communication system. The specs were submitted for straw ballot.

Kiel: The OIF's charter is much larger in scope than the ODSI. The Optical Domain Service Interconnect (ODSI) was launched in January 2000. It is the informal coalition of more than 100 companies working to define an open interface to the switched optical network that allows user devices (for example IP routers, ATM switches, and SONET/SDH add/drop multiplexers) to dynamically request high-capacity bandwidth from the optical network. Significant progress was made on the Functional Specification Review during April 6 meeting in Chicago. The three other documents, MIB, Signaling and Service Discovery Specifications, are all revised and available at www.odsi-coalition.com. More than 100 companies and 170 individuals participated in this second meeting.

The ODSI Functional Specification update was successfully presented at the OIF's annual meeting on May 1 in Montreal. The third ODSI meeting will be held on July 25 in San Francisco. The plan is to begin initial interoperability testing in the third quarter of this year and to demonstrate multi-vendor interoperability by the fourth quarter of this year.

COOK Report: What is your assessment of ASON (Automatic Switched Optical Network Standards development and where does it fit into the ODSF and OIF work?

Kiel: G.ASON (Architecture for the Automatic Switched Optical Network (or ASON)) is an important emerging ITU-T standard in optical networking. Its scope is the Control and Management architecture for Automatic Switched Optical transport network, which enforces separation of control from client networks. It emphasizes provisioning, not restoration but it does provide some means of restoration. G.ASON architecture takes advantage of the existing (and future) sophisticated fault-detection mechanisms available in the physical layer of the network. This standard activity just started.

Backing up a bit, the ITU working groups have a much larger charter than the ODSI. The mission of ODSI is to drive a specific capability and work with the appropriate standards groups to have it included in documents with larger purposes. ODSI intends to bring its findings to the appropriate standards bodies for submission as a working item. The intention is to drive a working implementation of a dynamic end to end lightpath through disparate domains. As happens with the ODSI, the OIF sends liaisons to the ITU for submission in the appropriate sub-working groups.

The ODSI and OIF are industrial coalitions or forums, while the ITU-T or ANSI are authorized standards bodies. Both ODSI and OIF can provide their specs to ITU-T via ANSI T1X1.5 as proposals for formal standard specs. So the related work (if any) in ODSI or OIF may be inputs of G.ASON (ITU-T is a worldwide standard organization).

COOK Report: In our July issue Noel Chiappa said: The one thing that is most likely to slow the commoditization of bandwidth is that the fundamental Internet architecture, you know, the protocols and the routing and all, isn't ready for it yet. It's all spit and bailing wire. Commoditization implies a tremendous amount of flexibility, a high rate of change in topology in the way things are connected together. It also requires a very robust infrastructure. That's one thing we certainly do not have. A lot of this stuff runs because there's a lot of smart people tweaking it all the time. And I'm not sure we have enough smart people."

What Chiappa said is a serious critique. How do the above standards work impact this issue? Is there a possible danger that the layer two and three issues might be worked out well and the whole effort run into a BGP4 brick wall?

Kiel: At the heart of bandwidth trading and

Continued on page 15

CEO of AIG Telecomm Emphasizes Difficulty of Commoditizing Bandwidth

Reaction to Commoditization Issue Depends on Amount of Dark Fiber Owned and Affect on Control Over One's Customers According to Eric Raab

Editor's note: Eric *Raab* has been the CEO of AIG Telecommunications LLC (AIG Telecom) since August of 1997. Before AIG, he was the Chief Technology Officer at IDT Corporation. He did advanced technology work at AT&T, Bell Labs in Murray Hill, New Jersey until he joined IDT in April of 1995. We interviewed Eric on March 31 both to ascertain his thoughts on the movement toward bandwidth commoditization and learn more about how technology changes in general are affecting AIG.

COOK Report: How has the world affecting you and AIG changed since August '97? One of the things you're doing is reselling minutes on the voice network which is a sizeable business in itself. But I have the impression that resale of voice minutes on connection-oriented networks might be quite different than the ultimate issue of a commodities market for IP bandwidth on connectionless networks. So with that broad blanket invitation, tell our readers what you see as the situation.

Raab: In one sense circuit-switched voice and IP transport are very different. And in another sense, they're very similar. You're getting information from point A to point B. That's the purpose of both of those. The differences come into how they're treated. For circuit-switched voice, it's very easy to label every transaction. So a telephone call is from point A to point B and it lasts a certain period of time. And at a certain time of day. And that's basically all you need to know.

Issues in Bandwidth Pricing - Voice versus IP

COOK Report: Yes, from the old days, there's a huge, complex billing system, time measurement and billing system in place.

Raab: Right. Think about the complexity of that, despite the fact that, for as long as you and I talk here, there's going to be one record generated in the switch that ultimately will be billed. So one conversation equals one record and whether it's difficult or easy, it gets translated into a dollar value that can distribute the economics down the chain.

Now, IP, or any sort of packet-switched sys-

tem, every packet that flows is an individual transaction. The beauty of IP is that you don't have to define how to get from point A to point B. One packet can go one path and the very next packet can go on a different path.

COOK Report: And obviously, though, anything that's reasonably billable is going to be made up of a big clump of these transactions.

Raab: Not only a big clump of transactions but one where the service level can vary, as well as the size of the packet, as well as the content of the packet. All of those components and many others enter into what the cost of that packet should be.

COOK Report: Elaborate on that a little bit more, if you would.

Raab: Well, you could imagine that if you were downloading a big file, say, from an FTP site, during off-peak hours, you might initiate the download before you go to sleep. You would expect to have the finished CD when you wake up. You don't really care about the speed of the transmission and you don't care about the quality, as long as over the course of the evening, you get the whole file.

COOK Report: But if you are having a mission-critical telecom conference of high level executives designed for a meeting to decide an urgent agenda. . . .

Raab: Which is the only type of meetings they have. All those packets have to be lined up and received one after another in a proper timeframe. And they can't be dropped and they can't be lost and they can't be retransmitted. Consequently there's going to be a different cost associated with that IP transmission. There are going to be different Quality of Service expectations.

COOK Report: For a moment, I was wondering whether you were implying a particularly fine level of granularity in deciding the costs of transactions. But it doesn't sound like you are.

Raab: Well, that's not necessarily true. Because if you're a router, then you're going to be processing calls from video transmis-

sions as well as FTP downloads at the same time. It all comes into the big pipe and then the router is going to decide how to treat it based on the information that exists theoretically. It's not there in practice yet.

COOK Report: Right, because Quality of Service is still kind of a holy grail. It sounds like we're getting there, more than we were a year ago, but it's not yet set in stone, is it?

Raab: It's not set in stone. There's no distinct standard that everyone has accepted and employs.

But even assuming that did exist, then the router has to decide on a packet-by-packet basis what all these parameters are with respect to the information exchanged and generate a record of it. And are you going to generate a record for every single packet? Well, that's going to be a huge amount of information. And if you take the complexity of billing the telephone network and multiply that by three orders of magnitude, that's the complexity of billing the IP network.

COOK Report: Isn't the point you've made, though why we don't generally bill things, at least not in the telco way on the IP side of the fence? And why we would generally be talking about a slice of bandwidth that might be a DS3 from point A to point B that would be rented or bought for an hour, a day, a week, a month, for some period of time. The purchase would be here's the whole slice of bandwidth, it's yours.

Payment Only for Usage as Goal of Commoditization?

Raab: Well, that's the way things are done now. That's not the model that the IP bandwidth monetizers are trying to get at. What they're trying to get at is you pay for the bandwidth as you use it. You commoditize the bandwidth in that way, you reduce the transaction costs and enable rapid transaction capabilities so that you can buy and use bandwidth only as you need it and pay for it only as you need it.

COOK Report: Help me clarify again a little bit more, because I have talked to some of them and what I'm not clear on is the definition of 'as you need it'. It sounds like I'm hearing a point of view that is somewhat telephony-oriented in the sense that in that world one might think of smaller bites of bandwidth, of smaller sessions, of more immediate needs - we have a videoconference today. Let's figure out how we buy the bandwidth for just the videoconference.

Whereas it's my impression that they are saying that if you want to buy a DS3 or an OC3 or whatever and you want it for any kind of long term period your choice was an IRU that used to be for 25 years. Now it's going down to 5 years and, indeed, probably less than 5 years. Nevertheless, if you're talking about a big chunk of bandwidth, maybe an OC3, you'll have to buy it about 3 years at this point. But with the falling price of the OC3, you've got to wonder, I know what the wholesale price of an OC3 should be in April of the year 2000. What's it going to be in April 2001? But with the commoditization of this, instead of buying it on a 36-month basis, one might be able to buy it on a 3 month basis or a 30-day basis.

Raab: I think the point of all this is, you have to look at what does commoditization mean. Commoditization means that one guy's bandwidth is identical to another guy's bandwidth. And the point of that is that you can interchange them at any point in time and pay a price based on the market. And so the natural evolution of that should lead to the ability to very quickly switch in one guy's bandwidth, switch out another guy's bandwidth.

COOK Report: And with the proper technology, could be a matter of perhaps even minutes, but very quickly right now might be, what? A day or two or an hour or two?

Raab: Well, that's where we get to IP commoditization. IP bandwidth commoditization. The whole point of the IP part of it is that IP can be rerouted instantaneously and therefore you would like then the ability to buy and sell IP transmission on a real-time basis as you need it over the course of a particular transmission.

Raab: Yes. That's the value that commoditization brings to IP. Shortening the term cycle from a 25-year IRU to three months is somewhat revolutionary. But it is not necessarily commoditization.

COOK Report: Agreed that the shortening of the cycle is but the first of many steps.

Raab: Okay. In order to get from, right now they've gone like 80% of the way. But in order to get to the last 20%, they've got to

Attendees and Purpose of Bandwidth Meeting

According to a March 23 Dow Jones News Service story — A group of telecommunications companies is meeting Thursday [March 23] to open discussions on industry standards for bandwidth trading. Set up by CompTel, a Washington, D.C., trade association that represents the competitive telecommunications industry, the meeting will initiate discussion among companies that are or probably will be active in a nascent bandwidth trading market that could be worth billions or trillions of dollars a year. . . . Tom Gros, vice president of global bandwidth trading for Enron, has two items he thinks should be at the top of the meeting's agenda. The first is establishing an industry organization to set the standard terms and conditions for trading bandwidth as a commodity. The second is setting a North American benchmark for trading. . . . Dow Jones Newswires was able to obtain a list of attendees. Among the invitees are companies with experience trading energy commodities - Dynegy Inc. (DYN); El Paso Energy Corp. (EPG); Columbia Transmission Communications, a unit of Columbia Energy Group (CG); and Koch Industries, which has investments in energy and telecommunications. A surprise on the list is MCI WorldCom. Some industry sources have said the company opposes the idea of bandwidth trading under standard terms and conditions. . . . Also on the attendee list are telecom carriers Teleglobe (TGO) and Global Crossing Ltd. (GBLX); NTT America, a wholly owned subsidiary of Nippon Telegraph and Telephone Corp.; and Progress Telecom, a unit of the electric utility Florida Progress Corp. (FPC). Also on the list is recent IPO Universal Access Inc. (UAXS), a company that matches carriers that have excess capacity with carriers that need extra capacity. LighTrade, a start-up company setting up pooling points to allow bandwidth trading, was also invited. Commodity traders are represented by Sakura Dellscher, Amerex, Prebon Yamane, AIG Telecom and the New York Mercantile Exchange.

overcome this technical issue that I outlined at the beginning.

COOK Report: My understanding of it is that they are very upfront in stating that, to really commoditize, you need something that you can call a benchmark at a benchmark price and a uniform industry contract. Also as you've implied so far in our conversation, some legally definable, understandable means of making it a really fungible commodity. Right?

Raab: Right.

COOK Report: And where does that bring us? I have the impression there are different ideas out there.

Raab: There is a pretty broad spectrum of ideas. At the March 23 meeting there was one floated that had an IP-specific contract defined. For bandwidth. But it didn't work, because, for the reason I said, the technology to accomplish it isn't really there yet. And what makes sense to commoditize today is, if anything, the pure pipe.

COOK Report: Okay, how are they defining the pipe, then?

Raab: You know, a DS3.

COOK Report: Wait a minute, a DS3, though, is a specified amount of bandwidth, so you assume that pipe is lit, right?

Raab: It's lit, but it doesn't have to be doing IP.

COOK Report: Sure, it could be doing SNA

or something if somebody wanted to do that on it, or it could be doing whatever.

Commoditizing Bandwidth versus IP Bandwidth

Raab: Exactly. So I think what we've done in this part of the conversation is I think we've come to a recognition that there's a difference between commoditizing bandwidth and commoditizing IP bandwidth.

COOK Report: I would agree with that. I have the impression that the marketplace is tending toward more and more IP pretty dramatically, but I don't think I know enough to say what percentage of the marketplace is IP right now and what it'll be this time next year or three years from now. What's your feeling on that score?

Raab: The biggest part of the market today is switched minutes. IP, bandwidth that's being used for IP obviously is growing more quickly, but growing even more quickly is the available supply of pure, unlit fiber. That's the supply side. The demand is that, I guess, primarily IP-driven. Demand for other products like frame relay and ATM, even though it has already a substantial, entrenched base, it's not growing nearly as quickly as the IP market.

COOK Report: And frame relay and ATM are kind of level two kinds of things.

Raab: Yes, and corporations use them for all their data applications. An application can

be an extended Ethernet bridge. It can also be an extended Intranet. But to really benefit the market, the transactional difficulty of getting data from point A to point B has to go way down. I am talking about the whole cycle — right now, if I'm starting from scratch today and I say, you know what, Gordon? I want to connect up our homes and I want us to be able to exchange real time video.

COOK Report: Right, we need a fat bit pipe between us.

Raab: That's it. I've got to call a bunch of phone numbers and get lots of assistance. They've got to come and run more wires from my home and more wires from your home. And the whole cycle can take three to six months. That applies to us, but it applies also to the major carriers. If they don't have interconnectivity between them to begin with, they have to go through that whole cycle and a big part of that chunk is actually figuring out who's going to be providing the service and how much they're going to charge for it. They have to develop an RFQ and then get bids on the RFQ. That's a long sales cycle, right? That's a problem right there.

COOK Report: And your point is, the example could just as easily be one of the pharmaceutical companies ten miles up the road from me in New Jersey connecting up to its other main office in Connecticut, right? Same principle.

Raab: Exactly. And that part of the cycle has to be cut out. We must shrink what typically is a three to nine month cycle down to a two to six week cycle.

COOK Report: Let me ask you one question. If the companies want to do an Intranet kind of thing, they can go through the path that you mentioned. In theory, at least, could they also not go to someone like a UUNET or Sprint and say can you provision us a VPN on your Internet service from point A to point B? In theory, if the Internet service has the bandwidth available, they probably could do it more quickly, could they not?

Raab: If you're already connected and if you have the right equipment on your end and they have the right equipment on their end, then that could be done pretty quickly, yes.

COOK Report: And if you're not connected immediately, it's probably then 30 days or some number of days to get the right equipment, but it probably could be done somewhat more quickly than just through telephone company A to carrier B to local exchange carrier C.

Raab: Yes, but that's a big if.

COOK Report: I'm just looking for a gradation of possibilities.

Raab: I understand. But remember that all that traffic is going to flow over telephone company wires. An unavoidable consequence of ordering service is that you're going to be dealing with a phone company's provisioning of your local service. So that is part of the transaction.

COOK Report: And particularly so unless you just happen to be real close to the POP of the major fiber carrier and have some wireless means of getting it from the fiber carrier's POP to your corporate campus.

Raab: Yes. And we all know how likely that is under normal circumstances.

COOK Report: Although it's probably going to become more likely in the future, but I see your point.

Raab: So what will happen in the future is that there will be more available bandwidth to the corporate and similar locations so that we can transition from the situation now, with provisioning in three to nine months including the RFQ and bidding to something done at the click of a button.

COOK Report: Right, people are telling me that is ultimately possible and it's not going to take terribly long, perhaps 18 to 36 months to get from here to there. How does that fit in with what you see?

Raab: I think that's a perfectly valid prediction. And then what you might want to be able to do is to be able to buy the transport capacity from any one of a number of providers, on an auction basis. And that's where the commodity aspect comes in, the actual transit part.

COOK Report: Just in the same way that natural gas and electricity is commoditized, right?

Raab: Yes and no. If you want to buy electricity from a secondary provider, you can do that, but you have to sign up under a new contract with that guy and get rid of your old guy. What you'd like to be able to do is go to your refrigerator and, if you turn it on one way, you have electricity from one provider. If you turn on the other way, you have electricity from another provider.

COOK Report: And a little dial. UUNET's price from midnight to 6 am is X and Qwest from midnight to 6 am is Y.

Raab: Right. Look at gasoline. Gasoline is a commodity. And shopping, choosing between those commodities is no more difficult a matter than driving down Route 1. So you needed an analog for that. But the dif-

ference is that you can't go and make an offer to the gas station guy. If you pull into Exxon and a buck forty-five. You won't get a very friendly reception by offering him a buck thirty. And that's really what you'd like to be able to do for bandwidth. The thing that's keeping us, the thing that's getting between where we are now and where that utopian vision is technology and standardization. And that brings us to the meeting last week.

Competing Interests May Make Standardization Difficult

At the Washington DC meeting last week [Editor's Note: March 23 - See sidebar], they were not dealing with the technology part of it, but rather with the standardization part. The goal is that you come up with one standard to define bandwidth. You solve that part of the problem and then you've got to deal with the technology. If you can define and measure a standard then, at the very least, you can sell it as Brand X and know it's not going to make a difference who you're buying it from.

COOK Report: And you can probably begin to do that more or less immediately, but a bit more slowly and a bit more time-consumingly and a bit more clumsily than just on a purely automatic basis. I get the impression that a lot of them feel it's probably necessary to go through all that to get from where we are now to where we all want to be. What's your point of view?

Raab: I agree.

COOK Report: Is it going to be relatively easy to define these standards?

Raab: Theoretically, it's easy, but practice is hard because you have competing interests.

COOK Report: Help me understand what those competing interests are.

Raab: Well, think back to the days of the telecomm wars, between AT&T and MCI. AT&T's position was we're going to protect our margin at all costs. And MCI's position was, we're going to sacrifice margin to get market share. So AT&T tried to differentiate itself for many years by promoting itself as the quality network. And MCI would come and say, We're just as good and, by the way, we're cheaper.

And that lasted for a while, but eventually the whole thing collapsed. And now, basically, if you can figure out how much you're paying for phone service — they try to make it so complicated that you can't — but if

you can, then you'll come to the conclusion that the services basically cost the same. Long distance has in effect become commoditized.

Now, same deal with bandwidth. All these guys — Enron, Williams, MCI — they don't want to commoditize their bandwidth. What they want to do is claim to provide a different, better, higher margin level of service than the other guy.

COOK Report: So you're saying that some of the participants in the Washington meeting, were more interested in establishing the type of agreements or contracts that would be tantamount to Quality of Service kinds of things. Is this what you see as going on, rather than something that is truly plain vanilla, fungible and interchangeable?

Raab: I think the issue has to do with who owns the customer? Who wants those customers? And the players who are most aggressive in commoditizing the bandwidth are also the ones with the most capacity to fill.

COOK Report: May I take a guess and say, for example, Enron and Williams, where the other side of the fence would be MCI?

Raab: You know the industry better than I.

COOK Report: The way that I would divide that is the ones that are known much more as raw fiber providers and to that extent, I think Enron may be but that Williams is selling primarily lambdas. But if you wanted a rule of thumb for deciding where you would stand, from what I'm hearing, it would be based what part of your business is actually selling IP bandwidth or bandwidth that people can put IP on and what part of your business is selling a wavelength of light? Like, for example, Williams does with the Sycamore equipment now.

If your business is predominantly IP bandwidth in the sense that MCI's business would certainly be much more IP bandwidth based than Williams', then the issue of who owns the customer is going to be more of a concern to an MCI than it is to a Williams. Am I warm?

Raab: I don't disagree with that.

COOK Report: If you want to characterize it in some other way, feel free.

Raab: I think, I don't see a major difference. It all depends, again, on who the customer is. And depending on the product, you're going to have access to different customers.

COOK Report: Are there gradations of customers out there that I may not be adequately aware of?

Raab: There are definitely gradations; I'm sure you're aware of them. Carriers have different needs. Telephone carriers have different needs from Internet service providers. And such differences focus on both on the type of technology that they need, but also the type of customer service and handling and restoration times and reliability that they require.

COOK Report: But were not we talking about a situation where someone like an MCI which needs to be able to provision OC3's and OC12's and OC48's from point A to point B and to do it easily, quickly, cheaply, effectively and so on, could profit from a commoditized contract that somebody like a Williams would be very happy to sell?

But, on the other hand, are we now faced with a situation where some of the people might be saying, well, that's cool, MCI could buy in really big quantity from Williams, but how about an MCI customer? Does this mean an MCI customer can buy its bandwidth from Williams and sidestep MCI? Or some other big carrier, Sprint or AT&T?

Raab: Yes, I think that's exactly the point.

COOK Report: Is that the dilemma that they have?

Raab: It's not really a dilemma, it's pretty cut and dry. I think the guys that have the customers want to keep them and the guys who don't have customers want to commoditize bandwidth so that their network looks identical to the other guy's network and then it's strictly a price issue.

COOK Report: So the scope of your commodity market becomes a very interesting issue. A large dark fiber owner like Williams or Level 3 could say, that if we have a commodity, not only can we sell it at the level of OC48's and maybe even OC192's to the big carriers, but we can also sell it at the level of DS3's to customers of the big carriers. Is that part of the issue here?

Raab: Not yet.

COOK Report: Okay, well, help me understand, then, better where we are now and how the people are defining the issues.

Raab: Well, the key to understanding it I think is where the bottlenecks are. When you say DS3 or even DS1, the bottleneck is the last mile. Once you solve that bottleneck, you're not going to have five different guys with DS1 into your house. You'll have one guy who's willing to make the infrastructure investment to give you that sort of capacity. So you, as a consumer, at that level customer, are really not going to have the flexibility to buy bandwidth on an auction basis, because whoever controls the last mile

is the one who controls the whole thing — unless of course you allow separate charges - one from the infrastructure provider for the connection and a second for the bandwidth that flows over that connection. [Editor: He's also overlooking the possibility that wireless may change this.]

COOK Report: But let's say we're businesses, for example somebody that a Nextlink might deal with to obviate that last mile question. Then let's get back to the issue of whether the concern here is the question of how readily, if I'm a big carrier, am I being asked to participate in something that from the git-go will make it easy for my customers to go buy bandwidth directly from the wholesaler and obviate me as the middle man? It sounds to me like this must be a really interesting issue.

Raab: Yes, indeed. Now, that's good for the consumer, not necessarily good for the carrier who currently owns those customers.

COOK Report: I gather that these were some of the issues that were on the floor in Washington bandwidth meeting. What can you say about how they seem to be shaking out and what people's positions are and where you think it's going to go over the next three months, six months, nine months. A source who was at that meeting claims that you seemed unhappy and left early.

Raab: Sure, but first let me say one thing for the record. I did not leave that meeting because I was unhappy with the proceedings. I left because I had another appointment. I understand people will say what they say, but that's the truth. And I actually returned to the meeting a little later on, but it had already disbanded.

But aside from that issue, that meeting was really somewhat of a misrepresentation by the organizers. It's not really their fault, I don't blame them, but it was really pushed by a few of the providers who had been trying for some time to promote their particular business model and standards into the marketplace.

An Enron Model?

COOK Report: Well, if I may say, from having talked to Stan Hanks and Lynn Franks, I have the impression that Lynn Franks' whole role at Arthur Andersen is to go out and talk about a model. But to talk about the Hanks, Franks, Enron model (or whatever the appropriate label is for it) and to in effect do some proselytizing of the model to the carriers to help them understand why embracing it would be to their advantage. That's my perception of what's going on there. Of course since Lynn Franks has worked with Stan Hanks while both were at Enron, one

might wonder whether in her Anderson work there is an agenda that is partial to companies like Enron and Williams that own a great deal of dark fiber and would benefit from bandwidth commoditization as a risk management tool? Does Andersen run a program like this on its own behalf without being paid by a second or third party?

Raab: You are asking reasonable questions. I think what Enron is trying to do is a good approach. I think that they had been encountering resistance among other players in the market because of competitive reasons.

COOK Report: I am just realizing that I did not ask about this. [Editor's Note: we did inquire about this within Andersen early in June and were told that such relationships were confidential.] And, as far as the other players are concerned, one may then ask who owns the customer? Is that accurate or it that not quite accurate?

Raab: That is accurate. I have heard it said among some of the people involved that Enron is now competing with everybody else, so no matter how good an idea Enron has, they're going to be reluctant to implement it.

COOK Report: Spell that out a little bit more clearly for me.

Raab: Very simply, just to take an example, I think it makes a lot of sense for two countries to agree on a border. No one denies that a border should be in place. Yet because of the nature of some countries, you will have a situation where they would seem to prefer to just slog it out for years on end, lobbing shells at each other. But at least in my limited view, the slogging looks like a useless and futile protracted war of attrition, simply because that's how they're used to competing, rather than working together. One must question the rationality of it all.

And not to stretch the analogy too much, but given the competitive landscape of telecommunications, in some sense you're going to resist the entreaties of another party, even if the idea is a good one and it's not completely clear that it's in everybody's best interest...

COOK Report: But one of the parties with a set of ideas is Enron and probably Williams is of a similar mindset? And some of the other parties are presumably folks like MCI or the current carriers with lots of customers.

Raab: Being in the energy resource area Williams is already a sibling to Enron and people I talk with believe that it can more identify with what Enron is doing.

COOK Report: Enron has its set of ideas and is it your perception that Enron was using the meeting to put them forth amongst people from both the more carrier-oriented groups like MCI and some of the more raw bandwidth-oriented people like Williams?

Raab: Some people think so.

COOK Report: Are all the parties at the meeting are continuing to talk to each other?

Raab: Absolutely. I mean, look, we're all in the same business. And we're all onto an exciting new frontier, which is the commoditization of network capacity. And the purpose of these get-togethers is really to ascertain whether there is a common ground for the development of that commoditization? And if that's the case, then we will work together. The only way to explore that is by having these meetings.

Summer 2000 - A Marketplace Overview

COOK Report: Okay, because some people say, (for example Enron) that on the basis of their own definitions, they're already beginning to do some trading.

Raab: Well, I'm certainly doing trading. So I'm sure that others fall in the same camp. I can't speak to the amount of business that others are doing, but AIG Telecom is certainly doing a trading business and it's definitely a viable business model. But we've done it by setting our own standards and enforcing our own rules and managing our own relations with our own customers.

COOK Report: And I get the impression that what is happening right now, and probably will continue to happen until standardization arrives, is that, if you want to buy it from Enron or you want to buy it from Williams or you want to buy it from MCI, you can call up and say what's your deal? And each of them will have their "deal." You compare the deals. Then you decide which one you're going to go with. Unless you just arbitrarily have a friend and you know he has a real deal on bandwidth this week. In that case you might well call up and do one-stop shopping.

Raab: Well, my goal as an operator of a telecom auction marketplace is do it better and more efficiently and more professionally than anybody else. And my goal is to capture the entire market by doing that sort of thing. And, as a business, that's what I've set out to do. If I am one of the first to do it and the first to do it properly, then I don't really care if ultimately there's nobody else who can provide a service that I provide.

COOK Report: And you're bringing us

around to the point of defining better what AIG Telecom is all about. It sounds like you're a middle man. You don't own great chunks of raw fiber, right?

Raab: We have zero fiber, we have zero network. We're not a carrier, we're not a provider. What we are is a marketplace where carriers can come and buy and sell a telephone minute on an auction basis. A real time auction basis. Where they have the confidence that they're dealing with an established company whose parent (AIG Insurance) has \$200 billion in assets on its balance sheet. They know we have a solid foundation.

COOK Report: Because we've just done a pretty interesting segue way into defining the business conditions that make your market model possible.

Raab: Correct.

COOK Report: If I have voice minutes to dispose of or in the future other kind of bandwidth, it doesn't make sense for me to have to do it all individually if I can go to someone like yourself. You can do it in such a way that, if you take a small percentage as your fee for services, it's going to be worthwhile to me because I save far more than that in the amount of time it would take me to do it on my own? Is that the rationale?

Raab: Right.

COOK Report: Okay, but if you're set up to trade voice minutes in a connection-oriented voice telecom market, in a general sense, there's some principles there that I'm sure are applicable to this other thing of trading IP bandwidth. But on the other hand, you will surely need to tweak your operation to make it fit this other area, that's close but still different than voice minutes.

Raab: It's different and it's similar.

COOK Report: Are you starting to do some IP bandwidth trading or just plain vanilla, layer 2 trading on which you can run various and sundry packetized networks on? What's your business model is at this point and where are you going?

Raab: At this point, as you correctly noted, we're trading primarily point to point minutes. We're also trading Layer 1, network capacity, on a more limited basis.

COOK Report: How do you define Layer 1 network capacity?

Raab: Physical layer.

COOK Report: Okay, but a defined amount of bandwidth?

Raab: Yes, like an STM-1 from New York to Paris.

COOK Report: And that's roughly equivalent to an OC3?

Raab: Yes. It's a good chunk of bandwidth. It's the thing that's the same, well, the thing that's different is that we don't deliver to just any point.

COOK Report: Well the major interconnect points at this point, Hudson Street in New York and various and sundry other places are known. Companies like Equinix are building bunches of what will become these interconnection points, right? The interconnect points are known, so can you just have a template up there that says, these are the interconnect points that we have relationships with? In other words, your marketplace may be defined in a fairly open-ended way. If I want some bandwidth, I can just come on and say, well, my unique situation is I have thus and such kind of need and, if your software is good enough, it could do some mixing and matching and partnering or connecting with potential suppliers on a very broadly different level?

Raab: It's not just a software thing, but also a hardware thing. In general however that is correct.

COOK Report: The hardware part of it is?

Raab: Well, you need some way to redirect the traffic from one provider to the other provider.

COOK Report: True. But I was thinking of it as being so entirely open-ended that the purchasers would take the responsibility. But if you're entirely open-ended and I came to you and said, my physical interconnect point that I would like someone with the bandwidth to meet me at is thus-and-such. Does it make sense to do it that open-ended? Or does it make more sense to say in the AIG marketplace we have three physical connect points or 33 physical connect points and you can pick and choose from any of these?

Raab: I think it has to be driven by either the buyer or the particular seller. It's not up to the auctioneer to identify the solution

COOK Report: Okay, that's more or less

what I had in mind. What do you see happening over the next three to six months to a year? What needs to happen that hasn't happened? Define the lay of the landscape and tell our reader what you think he needs to know to make intelligent decisions.

Raab: I think that the end user will start to see tremendous benefits and the carriers become more aggressive in making their service accessible quickly. Companies like AIG Telecom crop up and force the carriers, the providers and the users to come together more efficiently. And all of this together is driving some exciting new technology, which will also benefit everybody in allowing them to make maximal use of the available networks and to have natural control over how they use the network and how they pay.

COOK Report: As a final question what might the technology be that it's driving the creation of?

Raab: Technology that allows for IP billing. And technology that allows for dynamic routing and that allows you to rapidly switch from one provider to another.

In-addr.arpa -- a Technical Elaboration

Editor's Note: in-addr.arpa is a technically very important part of DNS that is seldom discussed. When a discussion did occur recently on BWG, we decided that it would be useful to reprint as a short article. We do so with the permission of Karl **Auerbach** and Milton **Mueller**. What follows is a description of in-addr.arpa and an explanation of how it works. Of course since BWG is an ICANN discussion list what follows, in addition to a technical explanation contains some speculation as to what ICANN my with to do with the tool.

Milton Mueller: (BWG list June 3, 2000) Would someone explain, slowly, this quirky little part of the DNS-IP address nexus?

I understand what's in the in-addr.arpa domain. I'm just not sure how and why it is used. It doesn't seem to be part of the actual resolution process. I thought the name servers contained resource records with the IP address.

Is the in-addr.arpa domain, rather than the zone files, where the world's name servers actually get those addresses?

Karl Auerbach: in-addr.arpa is used to lookup names from IP addresses.

It works as follows:

One expresses the name to be looked up as the backwards dotted notion of the IP address being looked up...

For instance, my main server, 192.203.17.71 would be "71.17.203.192.in-addr.arpa"

(Notice that the dots are at 8-bit boundaries. Since we went to CIDR this has thrown a bit of a monkey wrench into the in-addr.arpa mechanism.) Anyway, I hand that name to my resolver and say "get me a PTR record". DNS works as it always works and returns the PTR record, if any, that matches that name. A PTR record contains a character string containing the putative name of the host.

In my case, if you dig into my zone file for my 17.203.192.in-addr.arpa zone you will find a record of the form:

```
71      IN PTR  npax.cavebear.com.
```

Trouble is that this could be anything I want, even www.whitehouse.gov.

So good resolvers take it an extra step and hand the putative name back into the resolver and say get me the A records for this. Then a check is made to see whether the A records actually contain the original IP address.

Now you may ask, how did I get the 17.203.192.in-addr.arpa zone? Easy, I told ARIN. They have a registration record and a whois database just like NSI does. And they will update the DNS delegation records in their 203.192.in-addr.arpa zone.

The delegation works *exactly* as does delegation for SLDs inside of TLDs except that the nesting goes deeper.

Is all this important? You betcha.

Many sites, especially ftp servers and mail handlers, will check the source address on a connection to see whether it is coming from some address that properly reverse resolves. Frequently if it doesn't, the connection is punted. For instance, most mail that my machines toss is mail that arrived from an address that doesn't reverse resolve.

Mueller: OK, so it's a reverse DNS. Some more questions:

You (Karl) have an entry in the 203.192.in-addr.arpa zone because you "registered" it with ARIN. Suppose you didn't register it. What would happen? Who's responsible for registration? What are the sanctions if any for not doing so?

Airwire.net Displays Innovative Business Model for Delivering Broadband Wireless Service in Partnership with Florida East Coast Railroad

Editor's Note: Paul Lewis is President and Chief Operating Officer of Airwire.net, Inc. headquartered in Melbourne Florida. A graduate of Florida Institute of Technology, he has worked in technical sales positions for Digital and Sanyo. In the early 90s he was a reseller of large computer OEM systems. He then moved to ADC Telecommunications and their Kentrox (CSU/DSU) Division. While at ADC, he was promoted to help build their wireless local loop business. There he successfully introduced the industry's first fiber optic digital microcell — smaller and less expensive extensions to cell sites tied together by fiber. He joined Arraycomm to create the domestic business-marketing plan. Arraycomm makes smart array antennas. While Arraycomm did not perceive a great demand for local loop in the US, Adicom, another Silicon Valley startup, was willing to focus on domestic wireless local loop opportunities. He moved there prior to joining Airwire.net in January of this year. We interviewed him on June 1, 2000

COOK Report: So what prompted you to tackle yet another start-up?

Lewis: Adicom was taking a wireless voice approach to wireless local loop and, as an after thought, trying to provide some data over it.

COOK Report: Not what you wanted?

Lewis: Not really. It was essentially a G3 type of approach. The thinking here is, "Let's take a voice delivery system and try to enhance it to deliver data". This is the same approach that the traditional wireless infrastructure providers are taking with G3, which is going to cost huge amounts of money for spectrum and only offer 128 Kbps per second.

COOK Report: So this brought you to Airwire? Did it get started as a normal ISP? What caused it to go wireless?

Lewis: I came to Airwire.net because it was the first company I've seen that was providing broadband wireless data solutions consistent with my wireless local loop philosophy. Airwire.net was originally formed as Wireless Internet Services of Florida in March of 1998. The company was specifically created as a wireless ISP and used microwave technology to bypass the traditional phone company local loops. Prior to joining Airwire.net in January 2000, I was writing

a white paper on wireless data and local loop issues. During my research into the this market place I stumbled upon Airwire.net. I was extremely interested in their business model and immediately signed a Non Disclosure Agreement. I then began lengthy discussions with Don Turek (Co-founder, COB / CEO and a principal investor at Airwire.net). We eventually agreed that by combining our experience we could greatly enhance Airwire.net's existing business model. We are currently seeking investment capital, which will allow us to expand our business to a much broader client base throughout Florida and the U.S.

Developing the Business Model

COOK Report: How does your wireless business model and tower and interconnect deployment differ from that of more traditional approaches?

Lewis: In traditional cellular and PCS markets, "A" band or "B" band license carriers own the license for a particular market area. Consequently, they have total control of that market. In the traditional ISP marketplace every company, to some degree, has equal access to the same territory. When I first sold wireless local loop infrastructure, it was necessary to own or lease a license to offer wireless local loop data. Therefore it was necessary to convince the license holder, for that territory, to buy your particular technology. More and more often I heard the question, "Can you deploy in the unlicensed band because I cannot deploy in the licensed band?" Unlicensed spread spectrum offers an advantage as you can deploy a network without the high costs associated with operating in licensed bands. What you need a good business plan, a system that is robust, scalable and based on solid RF technology - and, of course, the money to implement it.

I wish that I could say that we had a methodical master plan in the beginning. Most of our early deployments were made under the duress of new business opportunities. This is common in our segment of the industry. I have found that very few companies in the wireless ISP business have invested in developing a proper RF deployment plan. There is a common misconception that wireless ISP networks do not need a lot of RF planning. This could not be any further from the truth. They are not unusually difficult for an experienced RF deploy-

ment person, however, they are not as easy as "I can see the site therefore it will work".

When I came to Airwire.net, I knew that we needed to evolve our business to include some very talented RF deployment experience. I hired Randy Ratliff (VP of Operations) who has over 20 years of experience at designing and deploying RF networks in virtually all RF frequencies. Randy has designed and implemented RF networks for entire countries including Jamaica and Kuwait. Airwire.net designs our networks to provide cellular like coverage. As a result when we advertise our services we know that we hooking up customers anywhere in our covered area will be cost effective.

We have now developed technology that allows Airwire.net to provide what we believe to be the most robust networks in our industry. We have several patents pending which will allow us to provide such things as access point site redundancy and fail over capability. What we are able to do is, instead of providing an access point with a couple of megabytes worth of bandwidth, is to tie together several 2.4 gigahertz radios to achieve up to eight times the capacity of a single radio. In addition we would get protection for the continued operation of the link should a single radio fail. We will be using these and other techniques to control the amount of bandwidth available to each PC on our customer networks. I can identify that a customer is doing an FTP and, as long as that customer is not sitting there all day long doing only FTPs, I can give him adequate bandwidth. While we use throttling capabilities of the radios, we also use their routing capabilities for the traffic engineering that we perform.

COOK Report: Wasn't Airwire.net already 2.4 gigahertz? Also, what do you consider the difference between being an ISP and an access provider?

Lewis: Yes, Airwire.net had already been using the 2.4 Ghz band. Traditional ISPs provide Internet access and related services over dialup. Airwire.net started as a broadband wireless ISP. I have modified the business plan to evolve the company to become an ASP (access service provider). An ASP can provide many types of services to many types of clients. Airwire.net also provides collocation services, fiber facilities (hard wired), VOIP (voice over internet protocol) services, and VPN (virtual private networks). We don't see ourselves as an AOL or a

Mindspring. We are more comfortable being known as a network service company that provides ALL (alternative local loop) connectivity to those that need it.

COOK Report: So, according to you, “access” is broadband, wireless, local loop connectivity to the Internet and in a more general sense, the most cost effective means of delivering IP data networking?

Lewis: It is even a little bit broader than that. We are talking about two distinct things. Yes, as far as data connectivity, the world is going to be IP based for connectivity and interoperability. Airwire.net uses IP based networks to provide broadband services to our clients. We will use a combination of wireless and fiber networks to accomplish this. On the other hand, we refer to “access” as a service offering. This means we provide our customers “access” to the World Wide Web and each other.

Accessing Fiber for the Wireless Local Loop

COOK Report: What is the ease of use of fiber along the Florida East Coast Railroad? Is it easier and less expensive than in most east coast cities? Are you laying new fiber or leasing it?

Lewis: We are not a fiber provisioning company and therefore we lease fiber from carriers. We are finding that railroads, highways and power companies have right of ways with abundant fiber and they are always looking for new innovative applications to help them grow their networks and subscriber base. Typically, alternative fiber providers don’t want to become phone companies where they would have to compete with established giants in high-density markets. There are plenty of CLECs that have established themselves in these areas and it is getting more difficult to add real value other than shear price. However, there is an opportunity for these alternative fiber providers to work with companies like Airwire.net where they can sell capacity and still keep their support costs relatively low.

We provide support to customers at connection rates that are below DS3s. Consequently the real synergy between Airwire.net and our fiber providers is that we relieve them from providing and supporting T-1 type connectivity. For example if someone calls EPIK Communications, the fiber selling division of the Florida East Coast Railway, and wants a broadband connectivity, EPIK would refer them to us. We can accomplish this in the most cost effective way, ether by fiber or wireless. Airwire.net partnered with EPIK during the final days of 1999. We will be in their co-location centers and have access to their fiber backbone throughout the state of Florida.

COOK Report: So exactly how do you interconnect? How do you bridge the “last mile?” Are you then ready to put a radio by the tracks where there’s no other link?

Lewis: Absolutely. EPIK is one of our partners. But we also plan to partner with CLECs and ISPs to provide them with an alternative backbone and delivery mechanism. Each partner has unique capabilities and when we all work together we can begin to really penetrate the “local loop” market.

COOK Report: Let’s imagine we are standing on that railroad right of way. Exactly what do you have to do to interconnect? What do your base stations look like?

Lewis: Our inter connection is a standard one Gig E (gigabyte Ethernet connection) that is handed off to us as an RJ45 type port. This is great for us as we simply plug that connector into our radios and broadcast TCP/IP data. We will soon be on net with EPIK in our 6000 square foot collation facility in Melbourne, Florida. This allows us to get our data in and out at speeds of well over 100 megabits per second throughout their network. Along the railroad right of way, we have 2 choices: We will either collocate in the EPIK POP or we would have to build our own POP. In either case, we have our radios located in a 19-inch bay interfaced to their high bandwidth Ethernet.

In most of the major towns, Boca Raton, West Palm Beach, Daytona, Melbourne and so on, they will have a POP of their own where they will break out fiber and bring that fiber to a building where we can co-locate. While we have rights to be in all their facilities, they also have what are called “cans” every 40 miles along the railroad right of way. At these positions EPIK will splice the fiber and provide us a fiber demark. We would build our own small facility by the “can” and they would splice in a one gig connection and leave us to execute the radio, antenna and tower. There will be some small inter lata areas outside the cities where we will need just this kind of radio connectivity because, for these unfortunate folks, everything is long distance. The kind of tower will be determined by the kind of radio propagation that we need at any particular site. We may see a water tower and decide to stick an antenna on top of that.

COOK Report: At how many points are you connected via towers along the rail line versus being in collocation sites?

Lewis: A couple. The agreement is very new and we will be building out on an “as needed” basis. Everyone is receptive and supportive.

COOK Report: Do you have a map of your current or your planned coverage areas including pops, co-lo, and radio locations that you’d let me publish?

Lewis: Sorry but I am afraid we consider that to be proprietary. Let’s say that we are going to bring affordable bandwidth to tier two markets by which I mean areas outside of the 50 largest US metropolitan markets.

COOK Report: So you are saying that you are going to come from the countryside and build up to the tier one city markets?

Importance of Good RF Design

Lewis: Absolutely. When I came to this company, I found that a lot of companies doing similar business plans had decided that they could take indoor LAN equipment go outside with it and create a business to support the link. But doing this takes traditional LAN people from a controlled environment into a hostile environment. You will find that under such conditions they will assume a lot of things, that for example, line of sight is true line of sight. But it’s not. You have to learn to deal with many RF anomalies. You can also increase power or work with antennas or some combination of these when radios don’t link. However people who have only worked with digital circuits are to “binary” in their attitude when confronted by many variables of RF design.

You can overdrive a signal it will go down. You can under drive a signal and it will go down. If you don’t have the Fresnel zone figured right, or there is swag (inversion), or a building is in the way, it won’t work. You have to be able to recognize and plan for these variables. You have to have good basic RF design at the heart of your network or you are not going to be able to provide a solid and robust service offering and if you can’t provide that, you can’t collect a premium for the service. I would be concerned about buying RF services from any company that does not have staff experienced in building and supporting these networks.

COOK Report: I gather from your emphasis on providing access services rather than being primarily an ISP is that if you have a customer who just wants LANs connected and wants to take care of his own Internet bandwidth, you are OK with doing only that?

Lewis: Yes. We have much flexibility in the types of services that we can offer and we have many choices regarding the kinds of radios, power, design, installation and so on. In almost every case there are bound to be multiple possible solutions including fiber hard wired solution. Now the LMDS model of a Nextlink is interesting. Except that I have never understood how to make it work. You have to cope with a fat, and limited distance kind of pipe.

COOK Report: And therefore you have to live in an urban area with high density to allow the customer to be within a mile or two or three at most to be able to gain access to a LMDS provider.

Lewis: Yes, in most cases the LMDS model is designed to interconnect between significant buildings in the urban core and then to concentrate on selling services off that network to clients in those buildings. Winstar, Telegent have tried this approach with poor financial results and now Nextlink is trying it too. Both LMDS and the new G3 technology are very puzzling to me. How do you pay outrageous money for the RF license and then build a cost effective network that fulfills the very high demand for affordable wireless broadband data to businesses and the public?

COOK Report: What is your opinion of the current G3 initiatives?

Lewis: Negative. Instead of following Qualcomm and Ericsson and the other big giants around, we are provisioning customers, at 128K and higher, every day. I have potential clients in the US and abroad calling me daily and asking how do we do what you are doing at Airwire.net?

COOK Report: How does the cost of your Internet access bandwidth per month compare with the cost of your competitors?

Lewis: We are competitive by coming in below the high priced Telcos and above the cheapest services. We win by providing the best quality and value in our service offerings. We do have a standard price list and begin there.

COOK Report: So when you are asked to price connections you will do so on a case-by-case basis and will undersell the "big boys"?

Lewis: Absolutely. For pipes above T1, we need to be able to provide custom quotes, which is typical of this type of service request. People will expect good clean data including local and long distance voice from one provider and at one "access" price. We offer easy to understand pricing and services while we commit to a minimum bandwidth data rate on a network that is extremely secure

Long Term Business Model

COOK Report: Given your market niche couldn't you take some of the hi-tech companies in the Cape Kennedy area of Florida and provide for most if not all of their telecom needs by total by pass of the local loop?

Lewis: Sure that is a possibility but to do that I have to build a network - in my case a data network. I do not want to become a CLEC; however, I can certainly hand off my customers to a CLEC that can complete their PSTN calls.

COOK Report: In addition to general telecommunications bandwidth, you are presumably selling access to the Internet. How do you connect?

Lewis: I am still signing agreements right now. I have in my NOC and collocation center a link that is expandable to OC192 coming from the MCI WorldCom pop, next door, in Melbourne. EPIK is also bringing in another a link that is expandable to OC192 and I am bringing in another couple of connections of which I can't disclose at this point. So Airwire.net will be not only be, multi-homed but will have links out to three other exchange points.

Nationally, as we grow, we will buy transit where we must and peer with, as many smaller players like ourselves as we can. We are presently expanding into Orlando and we have plans to be providing access to 80% of the popula-

tion of Florida by early next year. We will use EPIK and possibly others for our back haul, but, because we are using radios, we are far more flexible than other strictly wire line competitors. When we announce our service offerings into new areas it becomes a matter of taking the calls from people who want to be hooked up and not having to worry about BellSouth making the local circuits available on time. I know I can build quickly with the normal delays that more traditional approaches encounter.

As our build out becomes more dense, customers will find that their laptops will reach more and more of our pops from more and more places. We begin to look more and more like a Metricom without having to hang a receiver on every lamppost. In fact, as our penetration has grown we have embraced it and now we enjoy a truly mobile venue here in Melbourne. We have a Hummer that is wired as a repeater and it allows us to travel at T1 speeds while remaining on-line even above 70 mph. Also, we have the ability to take our laptops out of the vehicle and go into a restaurant and dine while being on-line. Our support techs love the fact that they can now go out and eat while still being "connected". Lastly, as we continue to grow, we will offer a very cost effective pricing model into the SOHO and residential marketplace.

COOK Report: How are you finding your staffing situation? How are people getting trained in the RF aspects of your business?

Lewis: The RF issues are fairly basic ones. Therefore we don't need to find extremely sophisticated CDMA experts. We can find a lot of folks with experience in corporate LANs and we will outsource a lot of our infrastructure build.

COOK Report: Is part of your business model getting the Florida infrastructure that we have been discussing up and running and then going to some VCs and saying here's what we have done in Florida. Here are the cost figures and the cash flows derivable and here's the investment that we need from you to start up in other areas. Is this what you have in mind?

Lewis: Yes for the most part it definitely is, except, I want and need to get the investment in now.

COOK Report: What radios are you using?

Lewis: When I came in January we were using primarily Breezecom and single access point radio sites. As I explained earlier we will be going largely to point to multi-point radios. Now, we also use Proxim and Wireless, Inc. radio equipment. Additionally, we are in discussions with other manufacturers. It is our job to stay on top of the industry, players and new innovations.

COOK Report: What can you tell me about your customer base including how many are ISPs?

Lewis: We have about 150 customers and 4,000 to 4500 end users. By customer I mean com-

mercial business. By end user I mean total number of employees of our customers who use our Internet access infrastructure capacity for connectivity. We have all types of clients from the Sheriff's department Emergency Operations Center, to the Medical Hope Mobile which is a mobile "free clinic" that provides free medical services to underprivileged people. While some customers are doing web hosting and related services, none currently are ISPs. We are an ISP ourselves with collocation offered; however, we do plan to partner with other ISP's. We have identified a couple of ISPs and CLECs who are disappointed at the way DSL has turned out for them and with whom we will talk about partnering or other kinds of "what if" capabilities.

COOK Report: So as a potential customer, I will decide what speed of service I want and for how many people and you will give me a quotation?

Lewis: Absolutely. Some customers want Internet for a single PC. Others want it for a network of fifty or more PCs. In some cases we need to provide only a single mailbox because they will run their own mail servers. In other cases we will provide 40, 50 or more mailboxes and help to monitor that side of the business for them. One of our customers is the School of Engineering at the Florida Institute of Technology. When FIT's wire line connection goes down, the Engineering School loans FIT its Airwire.net connection. Florida Today, a local newspaper subsidiary of USA Today gets its connectivity from us. The county Sheriff's Department has become a significant customer as well as other agencies of local government.

Sycamore Networks

Continued from page 6

bandwidth commoditization is a network that is very good at turning things on and off. There is work going on in several different areas that are focused on allowing networks to be more nimble. One of the first challenges is the existing underlying optical network or transport network. We are seeing that initial next-generation networks are first focusing on making that part of the network more flexible — the ability to provision, deprovision, and reprovision high-bandwidth paths on demand in milliseconds — essentially liquid light. The second area of industry focus is on more advanced IP traffic engineering capabilities — e.g., Diffserv and MPLS. As these technologies mature and are widely deployed, the network will begin this transition. This network architecture shift, in combination with more advanced traffic engineering capabilities with the IP network (e.g., MPLS), will allow commoditization of bandwidth to begin to occur. All applications of this new technology are being used within a single provider's network — extending this type of dynamicism across network and service provider boundaries is still very much a research issue.

ICANN Trying to Rig Elections to Maintain Tight Control

The globally influential news weekly, the *Economist*, published "The Consensus machine" in its June 10, 2000 issue. http://www.economist.com/editorial/freeforall/current/index_sf4096.html. It begins:

"THERE is nothing like an absence of regulation for stimulating innovation. Such was the creed of early Internet enthusiasts. As John Gilmore, a noted online activist, has put it: "The Net interprets censorship as damage and routes around it." The myth that the Internet has thrived only because it is anarchic is now firmly entrenched."

"Yet myth is what it is. In fact, cyberspace is highly organised and even regulated, and not just for technical standards. What is unique about the Internet is not that it is ungoverned; it is that its regulation has emerged from the bottom up and not the top down. "The Internet's true strength is that, as an institution, it exhibits characteristics of policy formation that appeal to one's sense of liberty," argues Joseph Reagle, a policy analyst at the World Wide Web Consortium (W3C), an Internet standards body," it concludes.

Co-opting the *Economist*

We find it hard to believe that this text was not written by a PR flak at ICANN's agency Olgivie and Mather. Every sentence is part of an artful formula. "Absence of regulation stimulates innovation." The favorite Clinton administration mantra. And to some extent true. The "myth that the Internet has thrived only because it is anarchic is now firmly entrenched." Ah ha! Gilmore's statement was a MYTH! The internet doesn't survive because it is anarchic. "In fact, cyberspace is highly organised and even regulated, and not just for technical standards." A hah. The internet is already organised and - how shocking - even REGULATED so those folk who accuse ICANN of wanting to CONTROL and REGULATE the net are wrong. The Internet is ALREADY controlled and regulated.

The *Economist* continues: "The process of policy formation on the Internet is not well known to the general public, or even to many avid Internet users. Besides the W3C, there is the Internet Engineering Task Force (IETF), which develops agreed technical standards, such as communications protocols, and its steering group, the IESG, which co-ordinates and approves them; and the Internet Corporation for Assigned Names and Numbers (ICANN), which oversees the system of domain-names such as .com and .org."

A hah! Vindicate for ICANN by associating

it with the IETF. Everyone knows and respects the IETF. Therefore ICANN can be given respect.

Economist: "These bodies have certain characteristics in common. They are largely self-created and self-governing. They are open in both membership and arguments, priding themselves on giving all voices a hearing. They are largely consensus-based in their decision-taking. And, so far at least, they have worked surprisingly smoothly, even surviving the wholesale commercialisation of the net. Yet the expansion and internationalisation of the Internet are likely now to put new strains on its entire bottom-up system of regulation."

What the spin meisters of ICANN want you to understand IETF works well. So does ICANN. IETF emphasizes openness and consensus based decision making. So does ICANN. At this point those of us who have observed ICANN's nearly 2 year campaign of DISINFORMATION see that it has worked. The *Economist* has been fooled. We gag in disbelief.

The *Economist* goes on to describe in detail the workings of the IETF letting ICANN bask in the glow. Then we read: "The remaining Internet ruling body, ICANN, has no choice but to become more than a club of citizen engineers. Its task sounds boringly technical: keeping track of the Internet's names and numbers. The organisation oversees the domain-name system, which links the 12-digit numbers that identify servers connected to the Internet to addresses such as www.economist.com."

"But ICANN's challenges are now more political and economic than technical. It represents all Internet users. Although the net is highly decentralised, its naming system is hierarchical, as it must be if every computer is to be easy to find. Only 13 "root servers" know where one computer has to go to find the address of another. Anybody who controls these has a life-or-death power over the Internet, says David Post of Temple University in Philadelphia."

Spin: ICANN has no choice but to be a little tougher than these doopy engineers. There is money involved. Dontcha see? The article goes on to admit that there has been controversy surrounding ICANN but goes on to excuse it all.

The article's final four paragraphs give the message of more ICANNs to come and ones that can really empower citizen control and democracy. The final lines announce that SOMETHING LIKE IT COULD HELP NARROW THE GAP BETTER RULER

AND RULED!! Dear readers we bring you the yoke of Orwell's Animal Farm and 1984. Wear it and be happy. Dear *COOK Report* subscribers, read them and judge for yourselves. Unlike most folk, if you have read our earlier ICANN coverage, you will know why we are so disgusted with what the *Economist* has published.

"Governments might learn from online decision-making too. The Internet's real promise for democracy may be less the much-ballyhooed electronic voting than the fact that the medium makes it easier for citizens to debate and inform themselves. The Berkman Centre for Internet and Society at the Harvard Law School, for example, is working on software tools to organise "deliberative polls" online. The German government plans such a poll for a new data-protection law."

"Internet-governance bodies also provide a lesson in transparency. They document everything and make it accessible online. The IETF e-mail archives allow browsers to discover why certain decisions were taken, even if they date back years. The W3C has an internal rule that nothing really exists that is not posted on the consortium's website. And ICANN posts transcripts of all board meetings, even telephone conferences."

"If ICANN succeeds in gaining legitimacy, it might one day spawn similar international organisations for other online policy issues with worldwide implications such as privacy (to avoid having a patchwork of different rules for the protection of personal information). That was actually the plan of Ira Magaziner, Bill Clinton's point man for the Internet until 1999, who wrote the white paper that first called for the creation of ICANN."

"It would be absurd to assume that politics could be solved Internet-style, but governments would still do well to study the online decision-making process carefully. Something like it could, perhaps, help to narrow the gap between rulers and ruled—one more example of how the Internet may have a profound effect on the offline world."

What a wonderful sales piece for the corporatist agenda we conclude. The corporate backers of ICANN have shown their hand. From the purveyors of the same deception (page 1042 of the *Loyola Law Review*) we get the following overview of Ira Magaziner's Keynote speech at the Loyola Law School Tech Law symposium held this spring.

"In many cases, however, the best type of legislative response to technology may well

be a very limited and circumscribed response. Reactionary or ill-considered laws may have the effect of frustrating technological advancement. In his keynote address that concludes this Symposium, Ira Magaziner, one of the chief architects of the Clinton administration's Internet policy, discusses this danger and advocates a wait-and-see approach to on-line regulation. Where collective action is necessary, Mr. Magaziner contends, the legislature may look to industry and other private sector stakeholder groups for self-regulatory initiatives. Such was the case recently where the U.S. government facilitated the private sector's regulation of the Internet's domain name system."

Elections - ICANN Style

So much for Magaziner's high flown theory. In this article, we present the continued duplicity of Esther Dyson, Mike Roberts, Louis Touton, Joe Sims, and ICANN in practice. In Cairo, faced with blistering criticism from CDT and other public interest groups, ICANN agreed to the public election of nine board members by September 30, 2000.

In mid May ICANN announced its rules for the up-coming elections. Michael Fromkin had the following to say about them on the ICANNWatch web site:

"Candidates will be chosen by a NomCom that has no one who looks like or represents the average Internet user. It also has no one from ICANN's 'loyal opposition' — the people who have been going to ICANN meetings or participating on line but tending to disagree with the existing Board. So that is a legitimate reason to worry. On the other hand, the NomCom promises "open and transparent" procedures including "posting its selection criteria, timetables, and updated procedures" online "for public review and comment". That's exactly what one would hope for, and were it not for what ICANN has done to the word "transparency" so far, you could breath easier.

Candidates will be allowed to attempt to run without the NomCom's blessing, but it may be difficult. The proposed rules for open nominations (ICANN calls it "self-nomination" although given how high they placed the bar "nomination by acclamation" would have been closer to the truth) have just been published today. There's a lot to worry about in these rules—they threaten to place unreasonable limits on outsider candidates, and also divide up the electorate, but it is only a draft so it may be too soon to panic. No, sorry, this is ICANN — it's time to post your comments and bite your nails.

Under the draft rules, in order to be nominated in the open process candidates must

get 10%, yes TEN PERCENT, of the eligible voters in their region to endorse their nomination. They must do this in 30 days, even though potential voters are only allowed to endorse ONE candidate. Since no one knows who the electorate is (voter lists are unpublished), and ICANN will be the conduit for a limited number of communications from the would-be candidates, you can imagine how many people are likely to qualify this way. And of course the bigger and more representative (and diverse) the electorate, the less likely a candidate can qualify. If four votes of the NomCom suffices to be on the ballot, why shouldn't, say, one hundred fellow petitioners suffice?

Other issues loom also: it seems as though the NomCom candidates might get as much a month's head start on campaigning, which hardly seems fair. Perhaps of greatest significance, however, it seems voters from each region will only vote for candidates from their own region. This is odd, since it is not how any other part of ICANN works. Members of the functional constituencies, such as the ASO or PSO, have to produce regionally diverse sets of Board Members, but everyone in the group votes on all candidates. Linguistic differences make this harder in the at large election, but that problem will exist even within regions.

One consequence of this geographic division is to make "slating" impossible — yet another way in which the existing Board keeps making suggestions that have the effect of entrenching existing majorities. Originally there were supposed to be nine directly elected at large board members. ICANN first tried to eliminate the direct elections, then when that caused protests, it reinstated them but cut the at-large contingent down to five, 'temporarily' until a date unspecified, so they wouldn't be too powerful. ICANN justified this on the grounds that the results of online world wide elections were too unpredictable (follow the link and search for "at large membership"). It would be a cheap shot to say that some people prefer unpredictable elections; the truth is that these elections are an experiment, there are many unknowns, but no more than there are for other parts of the ICANN structure. Furthermore, the strategy of dividing up the electorate will not scale up if ICANN ever does seat the nine at-large Board members who were part of the original design. There are only five regions, and nine doesn't fit well into five." Fromkin piece posted May 20, 2000.

On May 25 a lawyer member of the private BWG wrote: "As a non-profit corporation organized as a California Public Benefit Corporation, ICANN is governed by the California Nonprofit Public Benefit Corporation Law (Cal. Corp. Code Section 5110, et seq.). In these laws, there are actually detailed procedures for nomination and elec-

tion of Directors voted on by the members. However, The statutory rules vary, sometimes significantly, from what ICANN's "staff" has proposed.

Notably, ICANN proposes as a nomination threshold:

9. To obtain a place on the final ballot, an individual seeking self-nomination must meet the following conditions: Support from 10% of the At Large Members in her/his geographic region, as defined by the ICANN Bylaws; and Support from residents of at least two (2) countries.

But Section 5521(b) of the California Corporations Code provides: ... any person who is qualified to be elected to the board of directors of the corporation may be nominated: (a) By any method authorized by the bylaws, or if no method is set forth in the bylaws by any method authorized by the board. (b) By petition delivered to an officer of the corporation, signed within 11 months preceding the next time directors will be elected, by members representing the following number of votes: Number of Votes Eligible to be Cast for Director Disregarding any Provision for Cumulative Voting

| | |
|---------------------|--|
| Number of Votes | |
| Under 5,000 | 2 percent of voting power |
| 5,000 or more | one-twentieth of 1 percent of voting power but not less than 100, nor more than 500. |

See, <http://www.leginfo.ca.gov/cgi-bin/waisgate?WAISSdocIDp92516627+0+0+0&WAISSaction=retrieve>

There are materially different standards. Based on the most recent report of member sign ups, there were 7630 persons registered in North America. ICANN's threshold is 763, well above the cap. The California statute requires only 100.

Also, ICANN proposes that the self-nomination period close on August 20th, with the August 20th - September 20th devoted to campaigning, and voting to take place between September 20th and October 1st. Thirty days between close of nominations and beginning of elections. The California Corporations Code, however, provides: The corporation's articles or bylaws shall set a date for the close of nominations for the board. The date shall not be less than 50 nor more than 120 days before the day directors are to be elected. See, Section 5522(b). That's almost twice as long as ICANN is providing.

There's other stuff in there (like a requirement that voting taking in place in person at an annual meeting, with ballots for those who cannot attend) that ICANN is disregard-

ing. What am I missing? My assumption is that somewhere within the millions of dollars paid to Jones, Day to date by ICANN, one of their associates actually cracked the book on the Non-Profit Public Benefit Law. No?

On June 5, ICANN's deadline for comments, Bret Fausett, and Jonathan Weinberg, delivered blistering critiques to the ICANN web site and when ICANN's board met telephonically on June 6th, it found that ICANN legal staff had recommended election procedures that violated the California statutes under which they were incorporated.

On the sixth ICANN explained: "Whereas, the staff posted on 19 May a set of proposed rules for the self-nomination phase of the At Large election process; Whereas, a number of helpful public comments and constructive criticisms have been posted; and Whereas, the Board wishes to provide additional time for public review and comment and to defer final decision on the rules until its next meeting in Yokohama, while recognizing the staff's need for sufficient guidance to properly construct and implement the online interface and database components necessary to launch the self-nomination phase shortly after the Yokohama meetings; it is

RESOLVED [00.43] that the Board generally approves the staff's proposed rules insofar as they relate to the timing of the self-nomination phase, the data elements to be submitted by candidates for self-nomination, and the posting of candidate information on the ICANN website (Rules 1 - 5);

On June 7th an insider offered his interpretation:

The plan had been to adopt the rules, in toto, yesterday. Instead, they put off to Yokohama the questions of *how much* support candidates need to show in order to get on the ballot; whether at-large members can support the nomination of more than one candidate; whether the names of supporters will be made public; and what ICANN has to do to let the membership know that there are people trying to get on the ballot. Andrew's problem was that he hadn't been expecting heavy-duty challenges to be made to these provisions. When they came in (some of them at the very end of the comment period Monday), it was too late to deal with them adequately in time for the Tuesday morning meeting.

ICANN as Test Vehicle for the New Corporatism

On June 9th another BWG observer pointed out: see <http://www.icann.org/correspondence/sba-15may00.htm>

MIKE ROBERTS TO SBA: It would not be appropriate, in our opinion, for ICANN to adopt or observe a notice-and-comment model of regulatory decision making, [as it does not] offer the environment of community participation in deliberative discussions that characterize most aspects of the ICANN process.

The observer commented: This guy is a piece of work.

And on June 10th Milton Mueller said: It's irritating, but he's right. ICANN was created precisely to avoid the kind of procedural safeguards associated with governmental bureaucracies and regulation. The fact that it has quasi-governmental powers makes any sane person want to have those safeguards, but that's the basic contradiction built into ICANN.

As for the "community participation in deliberative discussions" stuff, well, is anyone here familiar with the theory of corporatism? I'm beginning to view ICANN and this whole international "self regulation" schtick as a variant of corporatism.

On June 11 Jay Fenello offered the following definitions of corporatism

<http://www.life.ca/subject/corporate.html>

Corporatism & Globalization — Increasingly, corporations are ruling the world. Many governments are, in effect, relinquishing control over their countries' economic and social agendas to globalization and the corporate agenda.

<http://www.britannica.com/bcom/eb/article/4/0,5716,26824+1,00.html>

Corporatism, also called CORPORATIVISM, Italian CORPORATIVISMO, the theory and practice of organizing the whole of society into "corporations" subordinate to the state. According to the theory, workers and employers would be organized into industrial and professional corporations serving as organs of political representation and controlling to a large extent the persons and activities within their jurisdiction. In actual practice, however, as the "corporate state" was put into effect in Fascist Italy between World Wars I and II, it reflected the will of the dictator rather than the adjusted interests of economic groups.

To this definition Mueller replied 'precisely what I meant'

We add that ICANN didn't take long to find a way out of its dilemma — it would change it's by laws to state that its members were not really members under the requirements of it's incorporating California statute:

On June 14 it posted; <http://www.icann.org/yokohama/atlargebylaws-topic.htm>

The Corporation shall not have members as defined in the CNPBCL (as defined herein), 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.

„(c) The Board may establish such temporary committees as it sees fit, with duties and responsibilities as set forth in the resolutions or charters adopted by the Board in establishing such committees.

On June 15 a knowledgeable observer had a very appropriate evaluation:

I find this particularly offensive. But it's so stupid I almost hope they go through with it. There is no action ICANN could take more calculated to demonstrate that they are not the body contemplated by the White Paper.

More and more, I've come to believe that the only political issue that really matters in the next six months is what Commerce does when the current contracts expire. If it just hands the legacy root to ICANN, then that's it - no more leverage. Commerce is, I believe, very keen to find an excuse to do this, although I cannot for the life of me understand why. It will, I fear, do so if there is any credible way that ICANN can be described as the body described in the White Paper.

And in the words of another observer: Under CNPBCL § 5056, you get statutory membership rights if you are "any person who, pursuant to a specific provision of a corporation's articles or bylaws, has the right to vote for the election of a director or directors." But the new bylaws, as drafted, don't by their terms give anybody the right to vote; they just state that five directors, to be seated this fall, will be "nominated and selected . . . according to a selection plan adopted by the Board." So it looks plausible we *still* aren't "members" within the meaning of sec. 5056.

So as always ICANN, publicly embarrassed, promises one thing and then delivers another. A continual shell game played with contempt for all those who stop to have a look at this wonderful new democratic formula by which, according to the *Economist*, the distance between rulers and ruled can be lessened.

Living and Working in Silicon Valley -- by Peter Deutsch

A Few Reflections from a Recent Silicon Valley Transplant

Editor's Note: On March 22 Peter Deutsch posted on the private Netscribes mail list an enlightening essay on what it was like for him as a Canadian early pioneer in the Internet to move from Canada and go to work for Cisco at its San Jose, California Silicon Valley headquarters. In 1990 -91 Peter was one of the creators of Archie at McGill University in Canada. Archie permitted the creation of Indexes for FTP sites and played a key role in enabling the web. Peter gave us permission to publish the following essay in a phone conversation during the second week of April.

A list member wrote: I'm struck by the intensity of the job market in the valley, and I'm meeting people on their jobs whose career to date was in politics or sociology.

Peter Deutsch: Yeah, I'm definitely seeing that and there's a real human interest story here. Here's a few reflections from a recent Silicon Valley transplant.

First, the work aspect:

I've seen articles in the Valley papers wit all kinds of horror stories, including thousands of open positions, right up through CEO ranks. Established companies such as Cisco are having big trouble competing with startups who offer fantastic leverage possibilities with pre-IPO stock options and lack of skilled staff are holding up legitimate opportunities. I'm fairly high up in my company, and am pretty happy with my package, but it's a funny feeling knowing I could get *more* if I bothered to shop around for a day or two.

In my group alone we have over 20 open requisitions for staff (something like a twenty percent staffing shortfall right now) and the company is paying a bounty to employees for successful referrals. Some positions even carry special "bonus bonus referrals" over and above the standard headhunter reward (know any developers or manufacturing engineers out there? :-). If you think this whole industry is growing fast now, imagine how fast it would all grow if we were fully staffed!

Many companies are looking overseas to bring in people, but the Immigration and Naturalization Service (INS) has now exhausted its supply of H-1 (work) visas for the year. There's talk of Congress extending the limit, but we're being told not to even consider overseas applicants this quarter, as they wouldn't be able to start for at least six

months.

The only exceptions are inter-company transfers and Canadian and Mexican workers eligible under NAFTA (that is, degree in a suitable field, work experience, etc). I've half-jokingly suggested sending press gangs across the border to do a little forced recruiting in the Montreal cappuchino bars. Recruiting seemed to like the idea, but we can't find enough large and beefy staff members to do the job!

Now the funny thing is that the positions being filled are actually very specialized and when you interview people you can tell very quickly if a person has the needed skill sets. Without them, they are of no use to you, as nobody has the time or inclination to train people up in a job.

I use to do that at Bunyip and watch new hires leave for a job that is "more fun" once they had time grade needed for the job at the next level up. Such workers are what I refer to as "gargoyle carvers", people who prefer to migrate from project to project, with their ties to their guild rather than their employer. And because of this trend, if you don't belong to an appropriate guild and match the needed skillsets, noone wants to wait for you to "grow into the job". You are paid well, don't get a real say in how the company is run, and are expected to move on to the next set of gargoyles down the street at some point. That's the new social compact between employer and employee for the techno-elite.

So, you end up with that infamous "two solitudes". Skilled technical workers have job insecurity, outrageous compensation and live in the overpriced housing on Doritos and adrenaline. Unskilled workers and Arts graduates make minimum compensation and are squeezed out the bottom of the stack. This cannot be healthy for society in the long run, but short of making sure everyone has computer science or engineering training, I don't have any concrete suggestions.

And that INS backlog has a human cost, as well. My kids are American, and I've applied for their passports, but we need a Green card for my wife. I'm told that the immigration backlog is such that it is expected to take as much as nine months before we can expect her to be able to legally enter the country to stay permanently, and her ability to enter on a temporary visitor's visa is entirely at the discretion of the immigration

officer at the port of entry, who must be convinced she really will return "home".

Because I can, for now I flyback to Montreal periodically, telecommute to work and otherwise work around the problem, but this is clearly not something that we want to see continue for too long. I'm grateful that I personally have the means to work around the problem but this clearly is not healthy for families in the long run, and for people a little farther down the stack this is not really a viable solution. \$100,000 doesn't go as far as it used to, you know.

And that's because, as you've all heard, the Silly Valley housing situation is absurd. I have a one room studio apartment (with kitchen and bath) in Mountain View that costs me \$US995/month. I looked around and that's a good deal. I pay a mortgage of \$Cdn720/month (less than \$US520/month) for a three bedroom starter home within 20 minutes of downtown Montreal. That same house would probably cost me about \$US500,000 to buy, or \$US2700/month to rent, anywhere in the valley.

My house in Montreal is nothing special, but then neither are most of the houses in the Valley I saw at half a million. We're going to rent and watch our stock options appreciate, but again most people don't have that choice.

Now, just about all the other costs (food, medical care, etc) are actually about equivalent to what I saw in Montreal (heating bills are actually down... :-). but that huge hole in your budget labeled "Housing" can't be ignored. Large numbers of "ordinary people" are moving away from the Valley because of the ridiculous housing costs, as the incoming techno-elite bid up the costs of living for everyone else.

So what does it all mean? It's probably kind of hard to feel too sorry for a 25 year old computer programmer making north of \$100,000 a year, plus options and bonuses, who complains about his or her working conditions, but the human side of this story isn't actually as much fun as you might imagine from a distance. There *are* perks, not the least of which is the knowledge that you are incredibly smart, surrounded by even smarter people, and at the top of society's hierarchy, but the numbers must be looked at in their overall context.

When France (my wife) blanched at the housing costs, I explained that we had to

imagine we were moving to the valley from Botswana. Of course the prices look silly compared to Botswanan prices, but we're not making Botswanan salary scales, either. Once you do the mental transformation, we're no longer rich in Botswanan or Canadian terms, we're just middle class Valley dwellers able to manage the payments here. It's a better life than many people here in the Valley, and better than the life we had in Botswana, but in this alternate reality plane, we're still only middle class. What's different is that I now realize that when I lived in Montreal, I was living in the technical third world.

Now, that may be a scary and possibly insulting thought for people living outside the Valley. Heck, it's unsettling for me, and I'm on the inside now. The challenge for each of us is to look inside ourselves to see if we're prepared to pay the price that the lifestyle demands, and then live with the answer, whatever it is.

The decision to move to the U.S. was a great stress on my family, and I know many people in Canada who could move who are not willing to pay the price. I admire them, but I now feel a wall between us when we meet. I can't complain about issues that are important to me without sounding ungrateful, so I tend to not talk about such things. At the

same time, when I hear someone with a Masters degree in computer science complain about the cost of living in Montreal my first reaction is "well, move to the Valley. You'll get \$120,000 a year plus options". This is clearly not what they want to hear, so I remain silent.

This is similar to the wall I noticed went up when I became the owner of my own company. I couldn't discuss the problems associated with firing staff at family gatherings when everyone there except me was an employee. I couldn't talk about the problems of financing half a million dollars in receivables after someone complained about the cost of swimming lessons. The gulf between us simply ruled out certain conversations.

And now I live in the valley, I see something similar in action when I come back to my old life. Other people's problems sometimes seem less intense. If I complain about my own I sound ungrateful. So I mix with my own kind and don't talk about my problems with those who I feel wouldn't understand them. And I suspect my friends outside the Valley think I've become more of a snob or at least more distant, as a consequence.

I'm not happy about this, but I don't know really what to do about it. And try as I might, I can't get that old shampoo commercial out of my head So I toss my hair and ask you to understand. Please - don't hate me because I'm a Silicon Valley techno-geek...

The radio ads seem to imply that you can get big bucks for breathing, and the cabbies are complaining that they aren't one of the 64 new millionaires a day. I had a CEO start a presentation with "You can tell you're in Silicon Valley when: * You call in rich instead of calling in sick * Nobody watches "how to be a millinaire" because nobody wants to lower his net work * Doctors and lawyers are the lowest paid workers..."

I saw one in the *Mercury News* a while ago that was almost on topic.

"You know you're in Silicon Valley when you see an ad for the new Palm Pilot with a naked woman draped across it, and everyone talks about the Palm Pilot..."

Okay, maybe it's not relevant, but as someone who's deeply involved with his new Palm V and trying to hide the true nature of this relationship from my partner, I can understand it in a way that I couldn't three months ago...

In-addr.arpa **Continued from page 12**

Auerbach: If I don't register with ARIN (or the appropriate RIR) then there is no DNS delegation for my branch of the IP address space. The result is that when somebody looks up any record for "71.17.203.192.in-addr.arpa" the query will fail just like a query for any other non-existent DNS name would fail. As for the responsibility - it is sort of fuzzy. Since I got my address blocks from Jon Postel, I try to expend my own efforts to make sure that I have things in order.

The registration could be done by the party doing the delegating or by the party receiving the delegation. Usually if an ISP is involved, the recipient does it. For delegations by ISP's to customers it is often the ISP that does it as part of the service. (We're only talking about the delegation here, the actual PTR records in the zone need to be maintained by whomever is at the bottom of the hierarchy.)

Bill Manning[*] apparently ran some scans and discovered that the in-addr.arpa structure was in better shape than most people thought.

([*] Apart from Bill's interactions with Chris Ambler, about which I don't want to express any opinion, I find Bill Manning to be doing a yeoman's job in evaluating the quality of DNS.)

As for sanctions - there are none - except for the indirect one caused by the fact that without a "working"[**] reverse mapping one may find that one's traffic is rejected by various kinds of servers, particularly mail and ftp servers.

([**] I put "working" in quotes because of what I noted in a previous message - simply resolving xx.xx.xx.xx.in-addr.arpa to a PTR record is only part of the way, but it is where a lot of checks stop. A full check takes the PTR record and tries to see if it can use that name to obtain an A record that contains the original IP address.)

Mueller: If the process of registration is as you say *exactly* like SLD registration then why aren't you paying \$35 a year for it?

Auerbach: I'd pay a fair fee for the service. Nobody has asked. The question is somewhat complicated by the fact that addresses are aggregated. And with CIDR the reverse lookup mechanism is far more cloudy and I can't say that I understand how the proposals work other than to know that they involve piles of CNAME (name alias) records.

Mueller: Should ICANN be paying ARIN for doing this, rather than ARIN paying 10% of ICANN's budget?

Auerbach: I'm sure that RIRs would like to be paid for it.

One could agree that just as ICANN is paying the IETF's protocol parameter registration costs by underwriting IANA (and at the same time ICANN is not paying the similar costs that are born by other standards bodies, such as the IEEE or ITU or W3C) - one could agree that maintenance of reverse delegation ought to be paid for by ICANN.

On the other hand, address space is a real resource, and I'd prefer to pay my fees directly to whoever it is that has to actually expend money to deal with my in-addr.arpa registration needs. I personally would not like to see it flow in to the pockets of Jones Day as it does when ICANN is involved.

Mueller: Any reason why ICANN needs to be in control of the .arpa TLD? What would happen if it was not? Are there possibilities for incompatibility/technical glitches?

Auerbach: ICANN does *not* need to be in control of in-addr.arpa except to the extent that it is the source of delegations to address registries. The sub-allocations are clearly in the domain (pun intended) of the RIRs.

By-the-way, it doesn't really bode well that ICANN went outside of the system and allocated address space to itself. I would suggest that it ought to consider forming a "provider independent" address registry. (This

Juniper versus Cisco

In Discussion on Inet-access list Juniper Rated Very Highly

Editor's Introduction: On June first on the inet-access mail list "Albert Fred" writing from a Hot mail account asked:

"When considering building a CORE NSP network, what are the pros and cons of Juniper and Cisco? Does anyone have comments that may sway our decision to choose one vs another? Who are some contenders to these vendors? We will probably use CISCO in the ACCESS and DISTRIBUTION networks."

The question by Fred produced a very useful discussion of the pros and cons of Cisco versus Juniper.

Phil Sykes: (Cable & Wireless, Europe) We're not deploying 12000s anywhere. Juniper have a proven line-rate infrastructure, stable software and a CLI that just makes more sense. They also have OC192 cards, great MPLS support in 4.0 and good techies on the phone. The just-announced internet processor 2 goes some way towards doing the things that the Cisco box does better (Netflow, ACLs). It's not a forklift, either.

Juniper downsides - lack of interface support (especially in Europe), although this is getting better, and some niggling 'why can't I do that's' in the configuration. We're not really even looking at Cisco for the aggregation routers either - the 7500 might be attractive again once the CzBus is finally done, and the C10K if the software ever works, but at the moment boxen like the redstone are looking way more useful.

Tom: I was told by a notable backbone provider who uses some distribution & core junipers that identifying smurf's on junipers has been a recent problem due to immature software that didn't have acl equivalents Juniper may have recently remedied this, but cisco's have nearly always been feature rich in this area with a proven track record & lots of documentation

Sykes: Not really a software issue - it was more an architecture decision. They didn't include packet sniffing/ACL capability on the original Internet processor because they couldn't do it without seriously hurting the boxes forwarding rates. They can do sampling of packets and ACLs in the new silicon, but I haven't checked out the features there in great detail.

Tom: Why did it take Juniper a few years to come out with anti-smurf tools?

Sykes: Not really a few years; more like 18 months. Juniper have a pragmatic engineering team who to a large extent are driven by feature requests from their big customers. We didn't need ACL capability, and I'll hazard nor did UUNET. Now we're using them for exchange/peering & distribution we do. Our core boxes are acting almost exclusively as MPLS switches; we'd no more focus our anti-smurf efforts here then we would on an ATM switch.

Tom: When considering the price, it may be a good idea to account for hidden costs such as expected downtime, support, product knowledge base, trained personnel base and vendor response time to security threats

Sykes: We certainly don't reboot our juniper boxen as often as our Ciscos. Support in Europe is apparently not a strong point, but the product knowledge base argument isn't a huge win for Cisco either - Juniper's documentation is a joy to work with compared to Cisco's, and the new version of the software promises to have explanations of all debug output - try getting that from CCO :-)

Chris Labatt-Simon: DISCLAIMER: We're a Juniper partner/reseller...

1) SHIPPING ASICs based wire-speed filter, sampling and deterministic load balancing card 2) Ability to filter on a wide variety of header information at wire speed 3) Interfaces up to true OC192, at wire speed (also DS3, SONET OC3 or OC12, ATM OC3 or OC12, OC48, GigE, FastE, E3, etc.)

Greg Ketell: (Juniper) And since the filtering is done in the forwarding ASIC in the "center" of the router you will get the same capabilities on any interface in the box as well as any new interfaces we come out with.

Labatt-Simon : 4) High density - 5 M20's to a rack, 320 DS3's to a rack, or using the CHOC-12, 960 DS-3 channels per rack

Ketell: Mix and match interface types on a single slot. ie 1 GbE, 4 OC3 POS, 1 OC12 ATM and 4 FE all on one FPC which takes up 1/4 of the M20, 1/8 of the M40.

Labatt-Simon: 5) Oversized ASICs - The M20 supports 20Gbps per unit, but the most interfaces you can fit in total 16Gbps.

Ketell: (This is what actually made me respond to the email. Sorry for all the other marketing stuff...) Uh, not quite. 3Gbps per FPC. 4 FPCs. 12 Gbps. Full duplex so 24 Gbps (counting as cisco does).

Labatt-Simon: The M40 supports 40Gbps per unit, and you can only fit 32Gbps worth of interfaces.

Ketell: 48 Gbps

Labatt-Simon: The M160 supports 160Gbps... you get the idea... 6) Single stage buffering - only one write and one read from shared memory across all interfaces.. 7) Software written by former Cisco and Bay engineers - non-legacy, ip only, ground up designed for speed and stability 8) Great CLI 9) Excellent TAC

With regards to why it's taken Juniper a couple of years to address things like smurf attacks... First of all, Juniper has only been doing general distribution of their routers since Q4 of 1999. Second, Juniper will not release hardware or software that degrades line performance. Processor based routers, even ones which have some cards containing higher performance ASICs, slow down substantially when you do filtering on higher bandwidth ingress and egress connections, and have quite a few limitations as to the filtering capabilities in general. Juniper waited to release their filtering code until they could run it at line rate.

Ketell: More basic than that. We released it when the ASICs could do it. The initial Internet Processor (the lookup engine) was the most complex ASIC ever built at the time. Technology progressed and we created the Internet Processor II ASIC which is pretty close to the most complex ASIC out there.

Combine that with you have to know what you are going to put into ASIC before you can... In JUNOS 3.2 we released the Route Engine Firewall feature which was basically ACLs protecting the router. We worked all the kinks out of the software filtering before committing it to silicon. The result? IP2 worked first time. Without proto-typing you end up spinning ASIC after ASIC as you work out the bugs. Just ask some of the Terabit router companies whose products still look really good on the slideware but exist nowhere else.

Labatt-Simon: The other item that most of our customers find surprising is that they are priced competitively against Cisco 120xx routers, and even similarly configured 75xx routers (but on the 75xx series, even with the RSP8's and VIP4's, I believe you are still limited to a 2.0Gbps per chassis).

On the other hand... no multi-service, a few missing interfaces, and there aren't thousands of VARs or direct engineers to support them - just a great direct sales and SE force, plus a few amazing partners like us <— shameless plug =) We still find a lot of people using them for regional aggregation and high speed edge requirements, in addition to core usage.

Ketell: It's only competitive if you look at price/port. If you add cost of power, cost of cooling, cost of rack space, cost of engineers trying to pick which version of code to run... Juniper wins every time. Then throw in cost/pps and/or cost/feature capability and there is no comparison.

Tom: thanks for the very informative synopsis of the impressive M160 internals using old ASIC technology. Your statements in the area of indirect costs: I wonder if yahoo or others, who have suffered a severe smurf attack, would agree that their problem was not significant enough to invest in the tools to diagnose the sources of those attacks.

Ketell: You are coming from history. We could not do it in the past. We can do it better than anyone else now. Internet Processor II ASIC can filter, log and sample all traffic running through it at any speed. The test was done with 32 OC48's at full wire rate (both directions) with input and output filters on every interface each with over 2700 lines of filters such that every line in every filter was hit by some packet being sent through the router. In this worst case scenario over 95% of full wire rate was achieved. Now you don't have to "turn on a feature to track and attack". You leave it running 100% of the time without fear.

Tom: I also wonder how cost effective it would have been from their viewpoint to have all of their connections through Juniper routers & then be told they have no hope with pinpointing the source in a timely manner. If this is not a problem now, then please forgive me since it's a moot point.

Picking a Cisco code release is only a problem on inet-access (heh). Competent engineers already know what IOS version to run. Cooling? Its been awhile since I visited a major nap, yet it was 89 degrees in there & everything that was plugged in, including junipers, ascend Cisco contributed.

Ketell: True enough. But which devices other than the Junipers were built to the customers specifications and then run @ 50 degrees Centigrade for 24 hours before being shipped to the customer? Which would be more likely to survive the heat?

Tom: Asics are a bit harder to upgrade than software when bugtrak publishes a firewall problem.

Ketell: You should really try to get in touch with a Juniper SE so he can explain how the system works. While you are right about the difficulty in upgrading ASICs, having to upgrade ours to a bugtrak problem is very unlikely.

Tom: Juniper's presence may be fueling the competitive process (a good thing in general), but "Juniper wins every time" just sounded too much like a sales pitch to let go unchallenged and all code (silicon or not) has a percentage of bugs that is unlikely to ever be zero.

Ketell: Never claimed it was zero. Our customers tell us it is dramatically better than what they were used to which makes it easier for them to scale their network. Isn't that what it is all about?

Tom: This isn't personal [but] the topic does seem to have a debate team sound to it.

Ketell: Debates are often best for enlightening the listeners. Participants may not change their mind but the listeners get to.

Tom: (If I were to suddenly change hats, I'd mention something like "Juniper never seems to mention 'code bloat'")

Ketell: We try not to bash publicly.

Tom: If Juniper is good on it's own merit, why do they have to mention the engineering team is composed of ex Cisco employees? (rhetorical)

Ketell: (Serious answer) the reality is that to be able to write protocols that work for Internet scale you have to have been writing protocols for Internet scale. Trying to take someone out of school who knows the theory and have them write your code just doesn't work. (Ask the Terabit companies.) Sooo Juniper swiped those who had already been doing it and gave them the opportunity to "write it the way you think it should be done". Juniper calls it the Internet Classroom. You have to be in the Internet to learn the quirks of trying to handle millions of users, hundreds of thousands of routes. We weren't there; we hired those who were. New startups have to try to do the same or they will not be able to have protocols that can handle the Internet well enough to get into the classroom to learn more.

Tom: Also, I was under the impression that Cisco code trains are customer driven.

Ketell: True. But they have far fewer Internet customer than Enterprise customers. So, how useful are those features to you.

Tom: Is FreeBSD really running inside the M160 router?

Ketell: Grey area. Back in 1996 the first development platform started out as FreeBSD. Since then just about everything has been completely rewritten. Way down deep surrounded by Juniper enhancements you can find what used to be (and still acts a lot like) the FreeBSD kernel.

Tom: Which vendor has more knobs?

Ketell: More? Cisco. More useful? That is for you guys (our mutual customers) to decide.

Dave Cooper: I think by asking whether or not GSR12000 is better or worse than Juniper M160 for running a backbone is prolly a question that cannot be easily answered (I am sure you can tell by the responses). To be honest both have their strengths and weaknesses and it takes some hardcore testing of features to determine what you should go with. Its like asking what ISP you think is the best, for every person that says ISP A is the best, theres another that says ISP A sucks. Best bet is, if you are serious and have the cash, get trial boxes from both vendors, beat the snot out of them and then you decide...

Ketell: Please! Test them side by side.

Cooper: a better question to ask the list is what test matrixes they used to come to their opinions on who is better and use the culmination of that input to build a better testing matrix.

Ketell: See <http://www.juniper.net/products/performance/performance.htm> this are the tests we use to try to best emulate the Internet. It was the basis for the DataComm tests, although they expanded it and actually made it tougher. Run any router through this series of tests and you will have a real good idea on how well it forwards packets in nearly Internet conditions. Then you will have to come up with your own Packet Filtering tests but a good start would be to run the above and then just turn on packet filtering on those interfaces.

Your own statistical results, coupled with third-party testing results will make your decision for you.

Coming in our September issue, out in late July, a just completed interview with Dave Schaeffer President and CEO of Cogentco which will be selling 100 megabits per second for \$1000 a month.

Executive Summary

Sycamore on Bandwidth pp. 1 - 6, 15

In the first of two more articles focusing in part on the commoditization of bandwidth we interview Jeff Kiel of Sycamore Networks. Kiel views a commodity market in bandwidth as a place where users will want to rapidly set up and tear down connections.

He points out that existing networks are not geared to respond to rapidly changing traffic patterns. In optical networks Kiel points out that the key differentiator is the software control of the light paths. Such software must permit one to switch and manage light paths. While Sycamore physical devices switch the circuits of light, Sycamore software pushes the intelligence of the network away from the idea of a central database and down as low as it can go, actually into the network nodes themselves.

Sycamore has implemented IP based routing inside its switches. Kiel believes that the real differentiation in optical networking is going to be not the size of the switch or the number of ports but how large can these networks actually scale. Could a network support hundreds of switches? Consequently the software within the network became the gating factor to everyone's ability to scale the data network. It must be decentralized because centralized management system can never scale fast enough so that it can deal appropriately with everything that's going on inside the network. But good software will also let the user choose the attributes of the circuit being established including issues like protection and restoration time.

Done right a portion of network bandwidth can be set aside as a protection capacity. With Sycamore's software it can be shifted almost instantaneously from place to place within the network on an as needed basis. Instead of the brute forces of SONET protection where half the fiber is set aside for fall over network utilization can be driven close to 100% by having a mix of protected, unprotected and preemptible circuits controlled by the network management software.

Routers are to be telling the optical switches to set up and tear down connections according to rules set by service level agreements and the need to partition the network into VPNs.

The Optical Domain Service Interconnect (OSDI) is an initiative begun by Sycamore in January 2000. The purpose is for routers to be able to tell the optical network to set up and tear down connections on demand. They want the IP layer to essentially call up optical layer bandwidth on demand. The role of OSDI is to get an initial set of specifications for doing this working.

To make commoditization a reality, what has to happen is that the underlying network needs to be re-architected and rebuilt to give that type of control to the carrier. And what the carrier will do is take that very nimble, very agile and very flexible network and use it as a tool to allow them then to move into the bandwidth commoditization and the bandwidth trading market. According to Kiel "There is work going on in several different areas that are focused on allowing networks to be more nimble. One of the first challenges is the existing underlying optical network or transport network. We are seeing that initial next-generation networks are first focusing on making that part of the network more flexible — the ability to provision, deprovision, and reprovision high-bandwidth paths on demand in milliseconds — essentially liquid light."

"This network architecture shift, in combination with more advanced traffic engineering capabilities with the IP network (e.g., MPLS), will allow commoditization of bandwidth to begin to occur." However the commoditization will be within the boundaries of a single network. Crossing provider boundaries is still a research issue.

AIG Telecom on Bandwidth, pp. 7- 12

In the second of two commoditization of bandwidth articles, we interview Eric Raab Chief Executive Officer of AIG Telecom. The discussion emphasizes the importance of understanding exactly what is being commoditized and of shortening provisioning times. Raab points out that the interests of all the players are not the same.

For example MCI could buy in really big quantity from Williams. But this could mean that an MCI customer can buy its bandwidth from Williams and sidestep MCI. Raab very reasonably suggests that carriers which have the customers want to keep them and the green field players like Williams and Enron which don't have customers want to commoditize bandwidth so that their network looks identical to the other carriers networks. Under such circumstances it becomes strictly a price issue.

Raab points out that the meeting of March 23 and two more that have taken place since then are designed to find out if the established carriers have enough interests in common with the owners of dark fiber like Enron and Williams to cooperate with each other in defining an industry benchmark for bandwidth trading. AIG itself operates a market place where carriers buy and sell minutes of connect time.

Airwire.net, pp. 13 - 15

We interview Paul Lewis president of Airwire.net a Florida based wireless ISP that is using 2.4 giga hertz spread spectrum two provide broadband wireless by pass to connect Florida businesses between West Palm Beach, Melbourne and Orlando. In addition to interconnection with MCIWorldcom, Airwire partners with EPIK which is the subsidiary of the Florida East Coast railway that sells access to the railroad's fiber.

Airwire is the first partnership we have seen between an alternative fiber provider such as a railroad or utility company and a wireless ISP dedicated to providing a local loop by pass for small businesses.

In -addr.arpa, pp. 12, 20, 24

A brief discussion between Milton Mueller and Karl Auerbach explains the operation of this obscure but vital part of the DNS system.

ICANN Elections, pp. 16 - 18

ICANN in fiction and fact — we observe how the *Economist* has been co-opted into writing an article that portrays ICANN as a regularly accepted part of internet governance — just like the IETF. The disinformation in the article reminds both of 1984 and Animal Farm. We chronicle ICANN's careful efforts to render the member elections to which it agreed in Cairo meaningless. ICANN has designed a process judged by legal observers to be in violation of the California Public Benefit Corporation Statute under which it is founded. No matter. ICANN proposed a by-law amendment stating that its member are not members according to the CPB statute. In the meantime ICANN refuses to answer questions about its 501 (c3) status. The June 26th *New York Times* has announced a lawsuit against ICANN by Afternic a reseller of domains which says ICANN violated its own bylaws in refusing to approve its application to become an accredited primary registrar for internet addresses." While ICANN claims concern about Afternic's reselling of domain names in the secondary market, ICANN's two largest accredited registrars Network solutions and Register.com "have begun offering resale services similar to Afternic's."

Silicon Valley Economy, pp. 19 -20

Peter Deutsh reflects on a move from Canada to San Jose where he is now working for Cisco .

Juniper M160 vs Cisco, pp. 21 -22

An Inet-access list discussion finds that Juniper hardware for its largest backbone routers is generally superior to Cisco and that Juniper's software is catching up fast. according to Juniper "with the new M160 and Internet Processor II ASIC can filter, log and sample all traffic running through it at any speed. The test was done with 32 OC48's at full wire rate (both directions) with input and output filters on every interface each with over 2700 lines of filters such that every line in every filter was hit by some packet being sent through the router. In this worst case scenario over 95% of full wire rate was achieved."

Continued from page 20

self-allocation is especially troublesome insofar as ICANN has removed its last vestige of legitimacy with regard to "technical coordination" when it removed the CTO position in the budget statement.)

Mueller: Am I correct in concluding that the actual process of delegating blocks of IP addresses is not affected in any way by management of (registration under) in-addr.arpa?

Auerbach: Right, the main concern of address delegation is to promote aggregation of blocks. The maintenance of the in-addr.arpa structure is simply a bit of back office book keeping that has no impact whatsoever on allocation questions.

Mueller: Suppose IP address delegation and management authority were divested from ICANN altogether. Would management of in-addr.arpa stay with ICANN, or be transferred to the new addressing organization? Put another way, is that .arpa zone more properly classed with address management or DNS management — or are the two so closely connected that Magaziner/Postel were correct in linking them together in the first place?

Auerbach: Well, to the extent that .arpa is a TLD, it would have an administrator/registr* just as does .nu or .mil or .edu. ICANN really needs to be involved only to the extent that .arpa is in its root zone. .arpa is pretty much free of a "character". It's the SLD in-addr.arpa that has the re-

verse lookup role.

And it is, in my mind, far more closely associated with, both procedurally and logically, with address allocation than with DNS. Especially if one considers that most of the DNS issues have been trademark oriented and it is hard to find a trademark in 71.17.203.192.in-addr.arpa. ;-)

Mueller: Does IPv6 have an equivalent of .arpa, will it use the same domain, or does it use some entirely different method?

Auerbach: IPv6 operates much like IPv4 but there are some new DNS records being proposed, at least in lieu of IPv4 A records.

Given that IPv6 is from the outset classless (unlike the Class A, B, C, ... of IPv4) and it is the residual relic that IPv4 space has been allocated on 8-bit boundaries (hence making it frequently possible to use the dots in IP addresses as equivalent to delegation boundaries) that makes in-addr.arpa work - the answer to your question is "I dunno".

COOK Report: Email will often be bounced back to the sender if the senders domain does not pass the in-addr.arpa reverse lookup test. True?

Auerbach: Yes, you are right. Most e-mail that bounces from the cavebear mail handlers is bounced because the IP address of the sender does not reverse resolve.

**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 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

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