The COOK Report on Internet



# LayerOne Gear in Telco Hotels Provides Cost Effective Optical Interconnect for Carriers

## Ciena Core Director Provides Service that Grooms Circuit Interconnection Between Fiber of Many Carriers

**Editor's Note**: Alexander Muse: is the President and CEO of LayerOne Inc. He has spent his entire career in telecommunications starting six years ago with ITT Communications, a competitive access provider based in Denver. From there he went to IXC Communications which is now called Broadwing. Next he ran a DSL company called Muse Communications. Before starting LayerOne he worked for Genuity. He began in the area of metropolitan networks, moved to the long haul area, and then, with Genuity, got into the IPworld. The common thread running through each of these businesses, a thread that he found made it difficult to do business, was the problem of interconnecting networks. He found that one of the issues in scaling, as a local access provider, was how such a company could efficiently connect to enough long-haul carriers to sell them services. We interviewed him on February 23, 2001.

***COOK Report***: I need to understand more accurately LayerOne's interconnection scheme. Would you for example be a customer of an Equinix Exchange point?

**Muse**: No. Equinix facilitates interconnection primarily at layer three with devices like Foundry gigabit Ethernet switches. We help carriers connect to one another. Equinix helps retail customers of those carriers connect to carriers. At Equinix, you run a fiber between yourself and your customer. They let you connect directly to your customer or supplier and charge a monthly fee for doing so.

You may go to a co-location facility like Equinix to buy access to a port. You may be buying private line or frame relay or ATM or Internet services. For whatever it is that you are buying, you simply pay the port fee. The co-lo facility also provides a distribution technology. For that service they charge you the so-called cross connect fee. They are adding value by making it easy for you to connect by means of an organized and efficient method. There are two pieces of your bill, the port fee and the cross connect fee.

In the case of Equinix as a co-location facility, because Equinix Centers are usually new and quite large buildings, they are normally not located at the areas of the greatest fiber density. Since such areas are normally down town (60 Hudson Street and 111 8th Ave in Manhattan for example), Equinix has to pay some carriers for bringing fiber out to their centers and making it available to their customers which are not interconnecting at layer one but are doing things at a higher levels in the protocol stack. To repeat: our business at LayerOne is connecting carriers together at layer one.

### Altering the Paradigm of Carrier Interconnection

***COOK Report***: Therefore because it takes less space to do what you are doing than it takes Equinix to do what it does, you are able to locate in existing 'co-los' or carrier hotels in core cities where the fiber density is high and where Equinix can't find enough room?

**Muse**: Yes. But the basic question should be what is their business model and what is ours? We are a carrier. We have "CLLI" codes assigned to us by Telcordia and are recognized as a point of interconnection between carriers. Every carrier site has a separate "CLLI" code. It is an identifier that enables carriers to do business with each other. We are recognized by other carriers as a carrier even though we do not carry traffic anywhere. But with the "CLLI" code we are able to accept and assign circuit facility assignments and we accept LOAs (Letters of Authorization)

**On the Inside**
**Issues in Metro Optical Networking**

to manage these circuits. In other words to turn them up and turn them off. We do circuit record locators and all the things you have to do as a carrier. Thus carrier interconnection is really our business model. I see Equinix's business model as marketing really nice, high security co-location space.

We can do what we do in as little as some 2000 square feet although our optimal space is about a 11,000 square feet. If we had 2000 square feet, we could likely only interconnect about 4000 strands of fiber as opposed to 31,000 strands with 11,000 square feet. The heart of our interconnection technology is our Ciena Core Director switch. It does not take that much room but the ADC Next Gen Fiber Frames and all the other fiber distribution gear surrounding the Core Director does occupy floor space.

Equinix is building new facilities. We are not. We are locating in existing carrier hotels and we are building inside these facilities what we call Nexus Optical Distribution Exchanges ™ . These Exchanges are engineered to facilitate the exchange of traffic between carriers at the most basic level, the transport layer – layer one of the seven level OSI protocol stack. We are offering the direct carrier interconnection services and related services to facilitate to the ability of carriers to take advantage of the opportunity to connect their fiber to our fabric and by doing so exchange bandwidth with each other.

We exist to alter the paradigm of interconnection. The paradigm in the past has been for a third party to build a ring architecture. That ring would interconnect the major buildings of the local market and then the architect of the ring would connect a SONET shelf into each of the long haul providers. This is done by companies like Brooks Fiber, ATT, MFS, and XO. The problem is that this architecture is not a cost-effective means of satisfying the bandwidth demands of today's markets. With a SONET OC 48 shelf powering a ring, every time you drop off a customer, you must subtract that customer's bandwidth from the total available. The problem is that you effectively strand bandwidth at each of these drop-off locations. But the only way to get around this is to install another ring, buy another shelf, drop off more customers and strand more bandwidth.

We find this to be a flawed model. It was, nonetheless, the model which dominated the early stages of the development of the Internet. You must remember that the NSF funded NAPs were OK for Internet activity until the total demand for bandwidth started to exceed the backplane capacity of the ATM switch. The problem now is that this architecture cannot support the connectivity needed for even the smallest players let alone the seven or eight large global transit backbones.

## Centralized Interconnection Better Than Gigabit Ethernet Shared Architecture

A telco co-location environment like 60 Hudson Street or 111 Eighth Avenue in New York City is typically where the long haul and the local access providers are going in order to do their interconnecting. A carrier will get 5,000 or 10,000 square feet and sometimes put in a raised floor with air-conditioning and generators-all the things that you think a carrier would typically do. Here in Dallas there is an equivalent building at 2323 Bryant with 56 carriers and over a hundred other providers co-located. This is a perfect facility for us today. Perhaps on the ninth floor you may have IXC (Broadwing) , on the 11th floor you will have WorldCom and on the 15th floor you will have Level 3. Now these providers buy services from one another. In order to do this they utilize the services of someone like XO. What XO then will do is bring its own fiber into the building and drop off a fiber shelf with SONET gear for interconnection on every floor.

But now the problem is it that if everyone must connect everyone else, and you do it on a one-off basis, you get a horrible mess. No one knows where everything is and we have a provisioning nightmare. When we have a tremendous number of pieces of equipment scattered everywhere, we find it that everything takes much longer to do.

***COOK Report***: In such an environment, how do you segregate the SONET players from the gigabit Ethernet World?

**Muse**: That is a good point. For example Telseon, who is one of our biggest customers, has said we are only Ethernet. So they buy dark fiber from MFN and today they're attempting to connect carriers together. But most carriers will not accept this as a solution because Ethernet is a highly shared medium and is by definition incapable of five nines performance since Ethernet packets will collide with each other when the shared medium is overloaded.

But rather than ask whether gigabit Ethernet is better than SONET or vice versa, I would have you ask how is the shared Ethernet technology being deployed? And the answer is exactly like the SONET technology is being deployed. That is: they buy a ring of fiber from MFN and they drop a Riverstone (a Cableton subsidiary) gigabit Ethernet box in each facility of each of their customers. They also have a box in our interconnect facility. They insert customers on what amounts to a gigabit Ethernet ring. What the Telseons of the world have done is created a big Ethernet pipe. In a year with the arrival of ten-gigabit switches, we will have a ten-gigabit pipe. And what is that? It is an OC-192. In the meantime, they have a one gigabit pipe into which they insert new clients or customers by adding a gigabit port card into their switch.

To sum up what we at LayerOne have learned it by looking at companies that must operate within this shared ring architecture, be it SONET or Ethernet or fiber in general, is that while there are interesting uses for it, it is not the way the you should be interconnecting your network to that of another carrier. It is also not the best way for a carriers ' carrier to inter connect and it is not a way for Internet service providers to interconnect with each other. It is not cost-efficient. It takes a long time and the architectures are invariably non scalable. What they've done is accomplished a one off connection with everyone with whom

they want to do business. When you had only four five major backbone carriers out there, this was not a big issue. But now if you include the carriers the CLECs, the DSPs and the ISPs, the total is more like 1700.

Consequently, in view of the optical technology that has now arrived, we think it's time to change the rules. We want to marry this optical technology with an architecture that enables carriers to do business, faster, cheaper and more reliably than they could before. A centralized metro inter connection is what you need not only for your internet networks but also for your long distance carriers and your ATM networks and your private line networks and so on. Moreover you need to do this not at the protocol level where everyone is interconnected today but rather you need to do it at the base layer. No matter whether you are Avi Friedman or Telseon or Yipes or Metropolitan Fiber Systems, the one thing that you have in common is that you are all running optical networks. See Figure One offering a block diagram of a Nexus Exchange on this page.

## Ubiquity of Optical Networks

What has become very apparent is that everyone is building optical networks. If you talk to *Telegeography*, they will tell you that there are approximately 1700 network service providers in the U.S. that use optical equipment. Forester claims there may be more than 6000. In our in Dallas facility we have over 45 carriers. We know of 150 carriers within relatively close proximity. While the new networks differ considerably in size from one another, they have many things in common and one of these is that intelligence within the network is moving toward the edge. Whatever protocols you are running, are running at the edge. And at the center of the networks at the first layer of the OSI we find fat optical cores made up primarily of the DWDM technology. Thus while we are using quite a few different technologies, we are doing so on top of a common architecture.

If you continue to examine these networks, the next thing you realize is that there is a proliferation of different network application technologies that ride at layer 2 and above. For example at Qwest, there are multiple networks riding atop the glass: frame relay, ATM, IP,

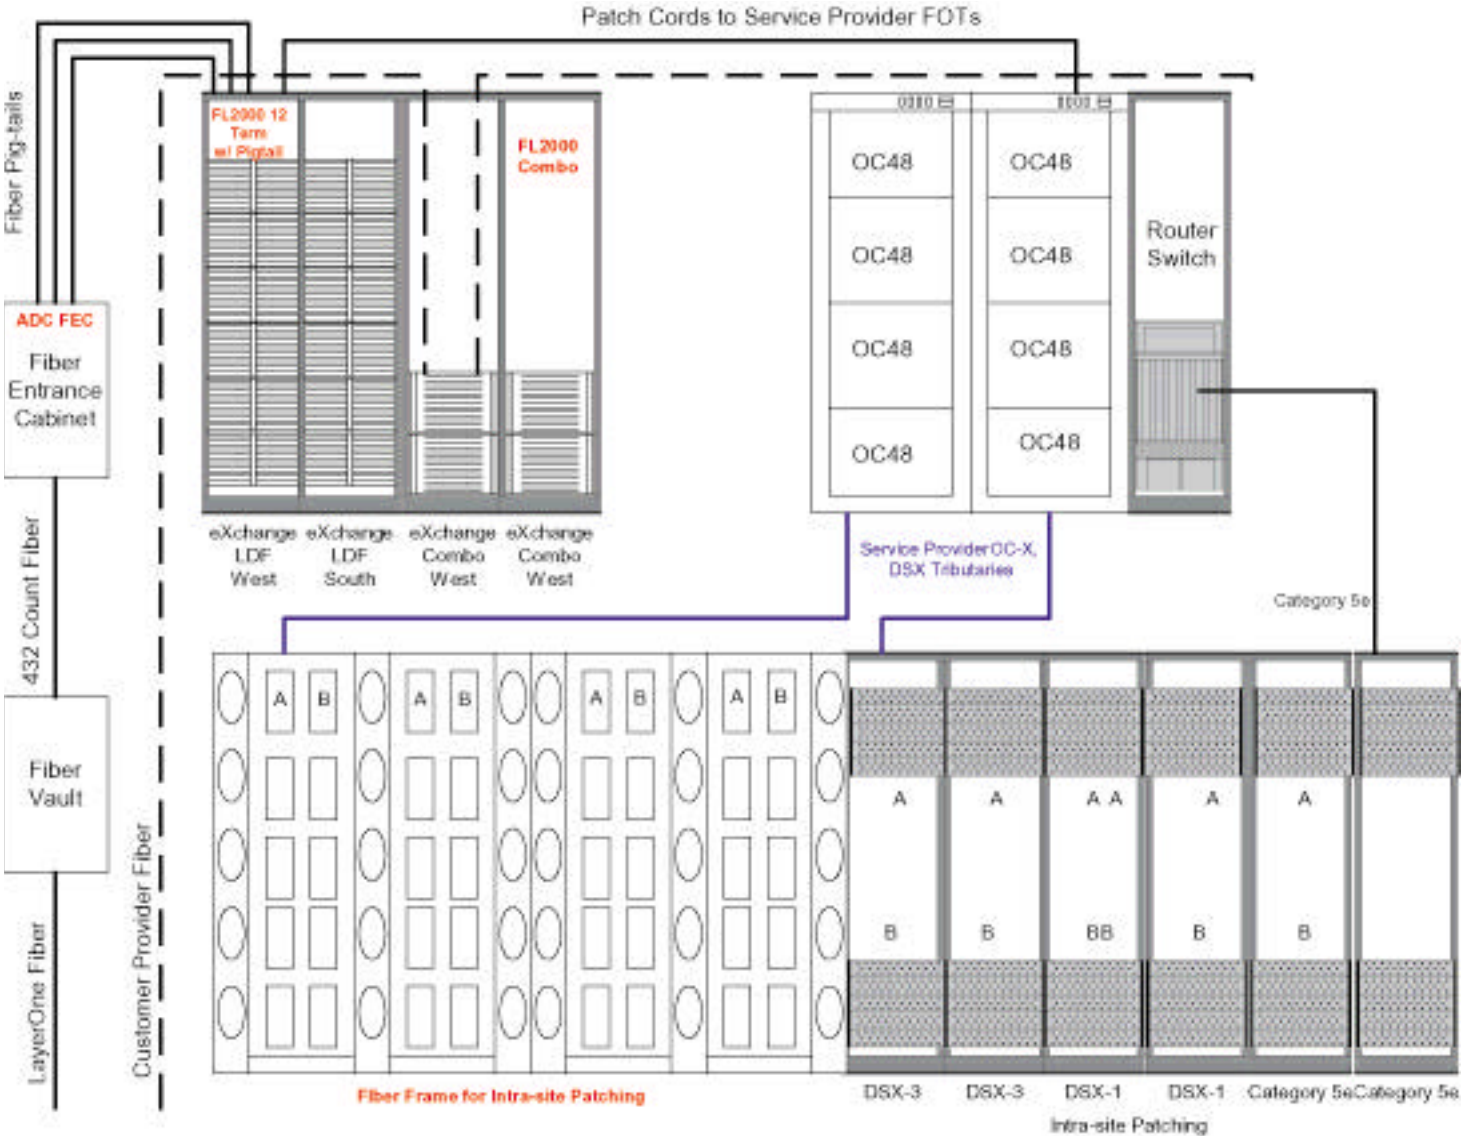


**Figure One: Equipment Used to Connect Carrier Fiber to the Core Director**

and voice long distance. You also start to realize that each of these providers see their applications as part of their own network and not part of some larger whole.

There are a lot of interconnections going on. The style has been to interconnect your carrier backbones based upon your network applications. These interconnections are carried out basically in every city. In any one city, at the edge of your network and behind your ATM switch and your router, they are provisioning connections to other folks. Also if you are an Internet service provider, you need to peer and/or get transit with other networks.

***COOK Report***: And you're doing it on a first-come first-served rather haphazard basis.

**Muse**: Indeed. You are doing it on a one-off basis with as many people as you can. Let's say you are a local exchange carrier. In that case, you are trying to sell local access to as many inter-exchange carriers as possible. You try to connect to as many people as you can, but you have to do this relative to each successive application at the edge. The problem all this presents you with it is that the interconnection of your networks happens at the worst possible place. At these spots each interconnection occupies a minimal amount of bandwidth compared to the total with which you must deal given the fragmentation of the industry where, instead of four or five major carriers, there may be up to 150 .

***COOK Report***: What you are saying is that for the cost of an ATM interconnect , you will get a lot less then you would if you simply tried to connect your bit stream with my bit stream at layer one?

**Muse**: Yes. If you look at these providers, you will see that what they do not have is the exact same set of matching applications. But that what they do have in common are fat optical backbones. Once they consider things from this point of view, they may come to the conclusion that they have an opportunity to take advantage of what they have in common with each other rather than what they do not.

While physically LayerOne does not sit in the middle of their networks, logically we have the opportunity to connect their own optronics and to take as a connection off the optronics where you currently run all the network applications that you have and interconnect that to a single fabric. If in every metro market these providers could connect to a single network fabric instead of directly to all the customers that they wanted to reach, they would in effect be connected from the best place in their network instead of from the worst.

***COOK Report***: Then this singular fabric is what you provide? It is the core of your business model?

## Enabling Direct Interconnection of Many Carriers

**Muse**: Yes. We have today the capability to enable a large number of carriers to interconnect with one another. Today the technology will enable you to do as many as 2500 OC 48 interconnections. Between these optically based carriers we insert what we call a logical core. This is a Ciena Core Director optical switch with a capacity of 256 OC 48 s. This switch gives us the capability of dividing that bandwidth in any combination of OC 3s all the way up to OC-192 interfaces. We can also go to granularities of DS3 circuits and one can expand the interconnected equipment up to a total of 10 bays worth of 256 OC 48 capacity bringing us beyond the promised 2500 OC 48 capacity at anyone location.

If you are a carrier and want to connect your various network application services to those of another carrier, you may do so as a customer of ours by connecting, to our logical core, the biggest possible pipe that you can give to us. We will then groom circuits according to your needs within that connection. For example if you are an ISPand you need to peer with 10 other circuit providers, we can enable you to do so within our core switching fabric. And we can enable you to do this with out having to go to an ATM platform or Gig E platform made up of a shared environment that adds buffering and error correction and is made up of PVCs and does a lot of interesting things. But we ask you: why have something adding that latency and overhead when you could have a direct optical interconnect to the network provider with whom you want to do business? You simply take a huge OC 48 or OC-192 pipe and give it to us. You let us carve it up for you into OC-12s, or OC -3s or even the smaller chunks of bandwidth, depending on what ever level of transit or peering you need to provide to the party with whom we are connecting you.

Now this is just talking about Internet interconnections. But this same carrier with a fat optical core likely used as voice and video bandwidth that might not be appropriate for an Internet interconnection. Because LayerOne is agnostic about the applications that travel on it, we can groom those circuits for interconnection with other non Internet voice and video circuits and just as easily as we can assist our carrier customers to do Internet peering. We are saying in effect why not interconnect optically with each other and by this gain the ability to layer on in each instance what ever protocol is desired by your bandwidth interconnection partners? If you don't want to use the IP protocol for your voice and video transfer, we will give you the ability to avoid doing that.

With the Internet we've designed a connectionless technology that is very good in having to deal with errors and recovering from them when they happen. But in the optical world at layer one, when we transmit just bits of light, we have far fewer errors to contend with. We can therefore safely interconnect far larger circuits. We think that in some cases, having shifted into the IP world, the potential for the economic handling of bandwidth may shift back into the time division multiplexed world. In my opinion this is why some of Avi Freedman's customers are telling him that they want traditional time division multiplexed voice circuits as well as IP circuits.

## Interconnection Architectures

***COOK Report***: I note that you are in Miami. Are you familiar with the NAP of the Americas?

**Muse**: Yes I am. It is right across the street from us. All the providers participating in the Nap of the Americas are also in our exchange. They get to the NAP of America's through local access providers which are in our exchange. Telseon connected to our exchange the other day through ATSI fiber. They are carrying a lot of data from providers in our state to the exchange. Then we have Cable and Wireless, Global Crossing, Level 3 and all those folk.

***COOK Report***: Do these folk do lower-level things with you across the street and higher level things at the NAP?

**Muse**: If you want to have someone sitting between you and Telseon when you carry out peering agreements, that is certainly something we're not going to do. So in that sense we're not all competitive with the NAPs. But a lot of the carriers who are really interested primarily in connectivity and can do their own BGP talk sessions and they have been initiating their own private peering agreements. In this sense we are a perfect opportunity when you would not want a third party to sit between you with an ATM switch or a gigabit switch.

In fact one of the limitations to our strategy is that some of carriers continue to do business in the way that they always have. That is to say sometimes people interconnect to our exchange from the edge of their network and not from the fat core of the center of their network. As a result, they only want to make a relatively narrow connection through our central fabric to their customer. But we are agreeable if they want to do this because, for example, some providers might only have an IP network.

We think the value of our exchange centers on the fact that we don't prescribe the method of interconnection. A fat pipe is most cost-effective. But almost any pipe is acceptable. Any provider connected to our exchange is welcome to use in the common software platform that all are given to determine who else is connected at what level. Such a provider by using the software would be able to determine that, for example, Cable and Wireless was connected with an OC 48 available while in the meantime I need to get an OC 3 connected to them. Our logical core exchange fabric also does a very efficient job of translating from SONET into a gigabit Ethernet. Our common denominator is that everything is operating optically.

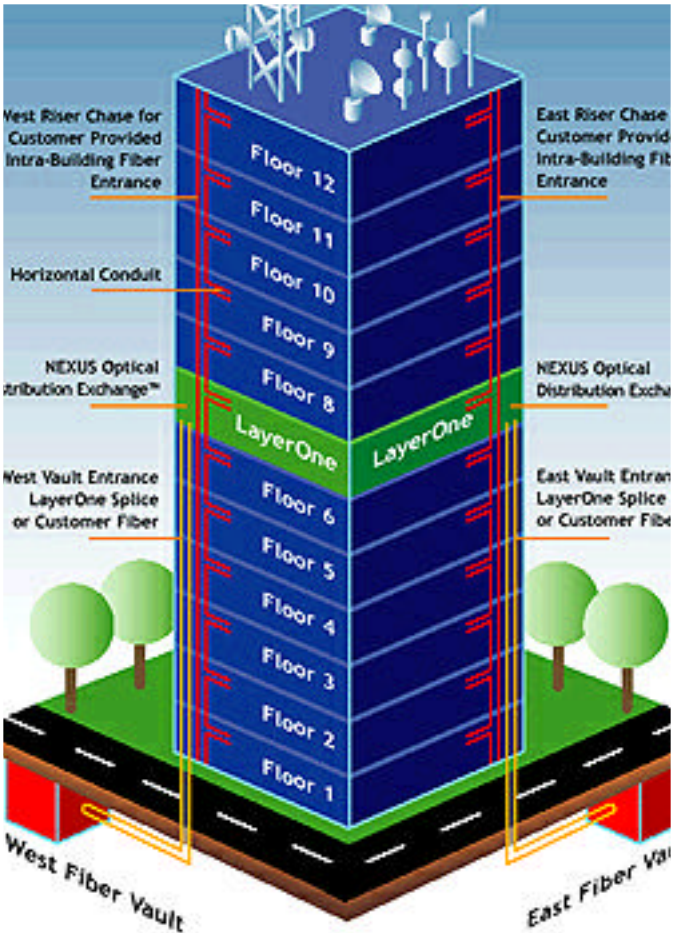


**Figure One: LayerOne's Position in a Telco Hotel**

***COOK Report***: So you are changing the dynamic from the idea of someone in a carrier hotel running a physical fiber cross connect from their equipment rack to someone else's rack when two entities want to interconnect to the idea of a single connection to your "logical core." Using your equipment and bandwidth management software, a carrier customer would configure one or many connections from his single interconnection to your " core."

**Muse**: This is precisely our point. You can all interconnect at one logical place or build your own mesh of physically separate connections.

***COOK Report***: Five or six years ago when the MAEs were just beginning this higher level protocol interconnection scaled. Now with the increased number of players in the marketplace, it no longer makes sense.

**Muse**: Indeed. With far fewer players and much less bandwidth, predominantly in DS 1s and DS-3s, it worked. Now with individual players needing very expensive optical interconnects such as OC-48s to multiple other players, our means of optically multiplexing our customers make economic sense. In order to do this we must change the basic architecture paradigm of how we think about interconnection to begin with.

We take a very large optical signal from you and convert it to electrical pulses in order to groom it for you into smaller circuits. And we can slice it and dice it for you usually in the precise manner that you desire. As a middleman specializing in this conversion for you and for many others, we can buy and operate equipment more efficiently than each of you can do for yourselves. In each metro market the cost to us of installing logical core conversion facility is 2.5 million dollars.

***COOK Report***: What does that 2.5 million dollars buy you?

**Muse**: The 2.5 million is basically for the electronics that does the conversions. The Core Director is a box about 7 ft. tall-roughly the size of refrigerator. This is the core fabric that will allow it to connect

seamlessly into everyone else. In addition to the core electronics, we need lots of fiber to be able to be easily brought into our exchanges by our customers. Therefore we have a framework of the ADC FEC boxes and the ADC FL2000 SC Connectors ELF bay. For electrical connections at the DS1 and DS3 level we use a customized ADC Entraprise Frame. Optical to optical connections at levels ranging from dark fiber to OC-192 we make via ADC's Next Generation Fiber Frame.

The equipment is shown in Figure One on page three above. It will support the exchange of about 31000 strands of fiber coming into our building from multiple sources. The sites that we have opened today have an average of about 7500 strands per site. We look to have growth into the 12,000 range over the next year so.

The 31,000 strands are there as the upper limit on what we think we will ever need to accommodate. The 7 ft. tall " refrigerator" or a group of them is what we would connect all that glass to.

***COOK Report***: In other words this rather complicated diagram is really a map of your ports or your entry ways?

**Muse**: Yes. That is the exchange. It is a combination of the active electronics and the physical devices that enable the fiber to get connected to those electronics. So let's say that Next Link (now called XO) would want to use our exchange. They would normally come in and bring their own opto-electronics — in other words bays full of expensive SONET equipment or Cisco gear. They would simply deliver us fiber strands — 96 or 192 strand cables. Initially they will connect two of them to the Cienna Core Director — each of which can handle up to 256 OC48s and each of which uses transponders and tunable lasers to do its work. Within the Core Director those would be then lit to a level of OC 48 or OC-192, depending on how they would want to interface. From there we would groom those to electrical STS -1s and then map those to the other providers to whom they wanted to connect. We charge them a flat monthly fee for connecting to our Core Director. The size of the fee depends on the bandwidth of the connection which can range from DS-1 to OC192.

***COOK Report***: Then you are taking optical equipment made by other people and integrating for your own purposes?

## Tributary Exchanges

**Muse**: If you talk to Ciena you will find out that they'd never imagined that the Core Director would be utilized in the way that we have. What we are saying is that there is an advantage to doing interconnection in the centralized place within a metro market rather than interconnection at the overlapping intersection of rings. We have deployed the product in such a way that we do not create rings with it but rather so that we create what we call tributary exchanges.

***COOK Report***: For example they imagined that MCI WorldCom would buy several of them for deployment in their own rings and that all the inputs to them as well as the out puts from them would belong to the single carrier owner?

**Muse**: That's correct. Other people have deployed them - including Broadwing and Genuity. They usually put them in the core of their networks. What we've indicated is that we are deploying the boxes without having a network of our own. What we are going to literally provide is the ability, should you have your own Core Director, for you to interconnect it with ours and insert the traffic that you want to bring into your network not from the edge of your network but from the core of your network. Now while you may not have a Core Director, you may have Lucent's Bandwidth Manager or Nortel gear and, if you do, you may use that for the same purpose.

It doesn't matter what gear you run. All you need is the appropriate high-density platform. Today the Core Director looks best but tomorrow it may be something else. We are married not to the technology, but to rather to the paradigm.

***COOK Report***: And to the bang for the buck.

## Version 2.0 Shared DWDM Ring

**Muse**: Agreed. Our current version 1.0 product does have some limitations. Our version 2.0 product-out late this year or early next-will be a shared DWDM ring. This should be very interesting for Bill St Arnaud and the other people at Canarie. The reason for this is that we are going to enable to the same density of connections to multiple providers and yet we're going to organize the connections in a different way by using DWDM technology.

Again this is a paradigm shift. DWDM technology is typically used between places like New York and Los Angeles. What you would do is to put a DWDM device on one end and a DWDM device on the other end and then you would basically put channels on that fiber in order to make it appear to be 20 fibers or 30 fibers or 60 fibers. We've determined that there is another application that you can have with this that would create rings within a single platform within a single exchange.

We will be creating shared lambdas and will be able to provide our customers who are connected to the fabric with interfaces to those lambdas. Moreover we will be able to do so relative to the application for which each of our customers is connected. For example we would have one a wavelength for BGP connections. Another wavelength might be used solely for OBGP connections and yet another wavelength dedicated to ATM connections. Or another for frame relay. We will organize the fabric by applications. Remember that I was concerned about the existence of a limitation that when an interface was available it might be only an Internet interface and not one that dedicated to a particular protocol.

With this environment if you wanted to send IP over any particular transport protocol, you would know that such a protocol would be available. Note also that this permits essentially no limitation on bandwidth on a single ring because you can concatenate or combine wavelengths. Therefore, theoretically you can have any number of connections to a sin-

gle ring from any number of providers and never run out of bandwidth relative to the ring.

***COOK Report***: In order to do this doesn't each connection have to represent a separate wavelength? Help me to understand how this is possible.

**Muse**: On your end you have interfaces that are relative to tributaries or potentially wavelengths. Typically, however, your interfaces will be associated with tributaries. A wavelength would hold a single tributary or multiple tributaries.

***COOK Report***: I thought a wavelength had to be an OC-192?

**Muse**: Not at all. The ITU has a grid and has defined different colors or frequencies at which photons of light can be multiplexed through a fiber. A wavelength is not relative to an OC 48 or an OC-192. What it is relative to is the number of OC 48 or OC-192s that can be interfaced onto it. You could actually have multiple OC-192 s on a single lambda or an OC 768 on a single lambda. The way you create a wavelength and push it down the pipe or the tributary is not important to the wavelength. The wavelength doesn't really care.

For a DS1 the underlying transport layer is LayerOne. Now you might put a layer two protocol like frame relay on top of it or even a layer three protocol like IP on top of that. Now in terms of grooming internally to the box, our customers always want to be able to groom down to DS3 levels. DS3 is an electrical level. But you are not doing a DS3 per se. You are doing what is called an STS -1.

## Cost Effective Centralization

In that environment you have been able to groom. So you may have an OC 48 connected box and you may want to buy a DS3 s worth of service from 10 other providers. We would groom that DS3 level out through the box. But the other providers who connect to the box would be connected to large optical pipes as well. In order to achieve this, what the carriers with their ring architecture today have to do is install a multiplexer into that architecture for every customer with whom they want to do business.

Of course, as you go out to buy the multiplexer, you have to guess how many cards to purchase for it unless you have assurance in advance of what customer wants. With our solution what you can do, instead of buying a multiplexer is to connect your fibers to our BIG multiplexer and from there permit us in to help you connect to anyone with whom you need to. We do this within a single device and giving you the flexibility to avoid stranding any bandwidth and to avoid buying something like an OC 48 shelf if you only need a small fraction of that much equipment.

Now come up with a single pair of fibers inserted into Core Director, and you can sell to 10 different providers or for that matter of a hundred different providers. When you do this you have no capital expense issues that you are forced to make. You have no difficulties in terms of provisioning time frames. You have no reliability issues of having to manage up to a hundred different multiplexers.

What we have shown is that, instead of distributing our connections, we can more cost effectively centralize them and do so relative to optical interconnects at LayerOne. In our version 2.0 product then we are giving customers the ability to operate protocol's in the their native mode over layer one technology. IP over wavelengths.

***COOK Report***: And the relationship of the version 2.0 product to version 1.0 is?

**Muse**: 2.0 is layered on top of our current version 1.0. Some of it is functionality that we will be able to add to the Core Director.

***COOK Report***: So where are you then in your national roll out?

## The National Roll Out

**Muse**: We have three exchanges currently open. Dallas Chicago, and Miami. We are under construction in Boston, Los Angeles, Minneapolis, Phoenix, San Francisco and St Louis. We have leases signed in Atlanta, Houston, Fort Worth, and a second Los Angeles facility.

Our current customers include : Enron Broadband, Verizon, FPL Fibernet,Giant Technologies, Genuity, Maxcess, MFS (Worldcom),XO (Nextlink), Progress Telecom, Qwest, SpheraOptical, Stonehenge Telecom, Wired Business, Yipes!, Cable & Wireless, Global Crossing, MFN, Tekom, McLeod Softswitch Level (3), Williams, Telseon, Cidera, Yahoo!, TXU, El Paso Global Networks, Epik, Edgix.

A master service agreement (MSA) means that a customer has agreed to locate in all our exchanges on a national basis. It sets up the parameters by which we will do business together on a national basis. We have signed MSAs with MFN, Williams, Verizon, Qwest, Level (3), Yahoo, Yipes!, Telseon, Sphera Optical, Cogent, Cidera, Edgix, EPIK, Innovative and TXU. MSAs are under negotiation with more than 40 other companies.

***COOK Report***: Doesn't Equinix have similar agreements with some of the same folk?

**Muse**: If you look at Equinix's S1, you will find some interesting insights. At one point we thought of collocating with them. In other words putting our distribution frames and Core Directors inside. One problem however is that they are located typically outside the central business district and not where the fiber networks are most dense. Consequently, what they've had to do is to pay a MFS, XO and GTE to bring fiber out their sites. If you looked at the S1 of Equinix, you will notice that Equinix has paid them $10 million worth of its stock to bring fiber to their exchanges. In addition they pay for monthly circuits. Equinix pays for this and it enables customers to co-locate there. Many of these customers are web hosting companies and dot com companies that use the fiber to buy access back to their Internet service providers.

Compare this to our Dallas exchange where we are interconnecting 29 carriers

ranging in size from Level 3, MFN, Global crossing, Qwest, Verizon and SBC to predominantly local entities like Global Metro. Numbers of strands of fiber range from 3500 to 16. Circuits lit range from 3 OC-192s, to 12 OC48s, to 76 OC 12s, to 90 OC 3s and more than 100 each of DS3s, DS1s and Ethernets. Total strands of fiber number about 16,000. We are contractually bound to tell the members of the exchange the identity of all the other members of and the number of strands belonging to each member. However we are prevented from giving this information to the outside world.

***COOK Report***: Who is Global Metro?

**Muse**: They are an example of a new phenomenon. Non carriers that buy fiber. There are real estate company that owns office buildings. There want to light up those office buildings so that they're more attractive to telecom clients. In thinking about the entire market for fiber you have to remember that it doesn't take the sale of many strands to put a substantial competitor in businesses. For example Qwest seeded Frontier and then through Frontier Global Crossing and also seeded Genuity by means of its original sale to the GTE . Level 3 purchased many of its strands from IXC, called Broadwing now and then sold one-fourth of them to Next Link now called XO.

Guys like MFN have built networks running hundreds of strands of fiber into buildings in outlying metropolitan areas. When you look at the Yipes! and Telseon business models, you find that to get into these buildings they only need a few strands. In this case however you must realize that your target market may be the entire fortune 1000. Let me tell you why. Consider a company like Electronic Data Systems (EDS) who had been spending $500,000 a month in the Dallas area on local loop fees — either with SBC or competitive local exchange carriers. These fees only get Compaq from their offices to the ports of the Interchange carriers. In a case such as this is potentially very attractive for EDS to buy say four strands of fiber from Metromedia fiber networks. If the strands are in the MFN city ring, they can get them say for a one time fee of $2 million or a lease of $20,000 a month.

In a case like this we have sales reps who work with sales reps from MFN. These guys go to a company like EDS on our behalf. While MFN can show the EDS how to put all its telecom and data services on those four strands of fiber, MFN is not very good in getting that data delivered to the proper carriers. That of course is where we come in. We tell EDS you buy a piece of optronics on your end and then chose to bring the four strands of fiber to us where we will interconnect them with our Core Director. We will then sell you interconnects with your dial tone provider, your long-distance provider, your high-speed Internet provider, as well as with any other kind of data on network service that you need.

***COOK Report***: In other words you can become a cost-effective middleman for someone like an EDS?

**Muse**: Indeed. EDS took a local loop monthly cost of one half million dollars came to us and to MFN a together we reduced it to a monthly cost of $50,000. As far as connection goes, Yipes! and Telseon did essentially the same thing. They took strands of fiber purchased from MFN and placing their gigabit Ethernet boxes in our Nexus Exchange, where they connect to other carriers.

At the end of the day MFN is rather like the guy who sits at the top of our food chain. They are really good to have. They help to ensure that a lot of people come to do business with us. But they are just one of the players. Level 3 is another. Level 3 has local dark fibers and long-haul. So does Qwest and so do a lot of other folks. Our exchanges ensure access to both metro and long-haul carriers.

Of course long haul and metro carriers exchange traffic but long haul carriers also exchange a lot of traffic with each other especially WorldCom and Broadwing. Both are customers of our Dallas exchange while in Miami basically all of WorldCom's Latin of American traffic goes through our site.

***COOK Report***: What are your criteria for choosing location of a new Nexus Exchange?

**Muse**: In every city we look for the carrier hotel or hotels with the highest density of fiber. In New York City we hope to be signing a lease at 60 Hudson within 90 days although a feasible alternative would be 111 Eighth Avenue.

## Evaluating the Market

***COOK Report***: What would happen if there were no space at a location where you really need to be?

**Muse**: You choose to do something as close as you possibly could. In Seattle for example there is a building that abuts the Westin Building which is the best building in all of Seattle for carrier interconnections. However there's no more room in the Westin building. What people did was go to the building that abuts it and blew a hole in the wall to let the carriers come through.

The closer in proximity you can get to the very best places, the easier it is to do your business. At the end of the day we could be anywhere and eventually everyone would build to us, *if* they liked our model. We went for best locations first and signed as many national agreements as we could until we got everyone in the habit of doing business with us. We made it as easy as possible for the largest number of carriers to do business with us quickly. Then, as we expand, we leverage the attractiveness of the early connections that we've made — in other words the attractiveness of our franchise. We should be able to convince them to come to some other locations that are not quite as convenient because they have seen the advantage of dealing with us.

***COOK Report***: If you look at your customers in Dallas how many of them have built national networks as opposed to having bought fiber from other people?

**Muse**: You raise an issue that is much misunderstood. In my opinion really only

AT&T and Sprint have laid national fiber networks. Qwest and some of the more recent players like Level 3 have certainly laid fiber, but have also wherever possible have bought existing fiber to fill in gaps.

***COOK Report***: So they do swaps. Let me rephrase my question. If you look at the fiber market what portion of it would seem to be occupied by owners of a lot of fiber infrastructure as opposed to what percent might be made up of newcomers just the buying into the market?

**Muse**: That would be a hard question to answer. I will tell you how I look at the world. If you own a strand of glass, it doesn't matter to me whether you leased it or whether you laid it or whether you have an IRU for it. At the end of the day that is your network and that's the only way I see it. And there are folks who are carriers'carriers who sell to other carriers and there are carriers to sell to the Fortune 1000. Most of the carriers out there do a bit of both. IXC is a classic carrier's carrier. Before perhaps 15 years ago you never heard of IXC if you were not in the telco business.

Their biggest customers were MCI, WorldCom, Sprint, Cable and Wireless and just a few other providers. They sold either dark fiber or lit fiber. They sold tributaries - perhaps OC 3s or DS3 s and they only sold those to other carriers. And those carriers made what they bought from IXC into their own network and may have resold much of it to others. Then they merged with Cincinnati Bell and were renamed Broadwing. As a result, they have a huge wholesale carrier to carrier business and a retail business focused on Cincinnati Bell — one that includes Internet services. Now Telseon may be thought of as an Ethernet player that only does wholesale services. They do not sell you Internet access. Instead they just offer a shared gigabit Ethernet pipe to someone else. They only sell to carriers to ISPs or to hosting companies .

No carrier has a truly ubiquitous network. Carriers talk in terms of 'their networks" while in reality they often grow or complete their networks by buying or swaping sections with some other carrier. In view of these fuzzy boundaries, you must realize that so much of what happens is really unrelated to current ownership. No matter what kind of attempts you have to brand services, they tend to remain a commodity because there are simply no distinct boundaries based on the ownership of infrastructure. While one carrier may outright own or have more long term control over a larger percentage of the fiber it uses than another carrier, there are no significant boundaries of long and short term ownership or control that affect day-to-day carrier operations. What they sell to the public and to businesses is a combination, not just of what they have built, but also of what they have borrowed, swapped, leased or resold.

***COOK Report***: So how would you begin to sum up your own market position?

## LayerOne's Market Position

**Muse**: We see ourselves today as having four major requirements for the success of what we're trying to do so. First, we are focusing on a highly interconnected venues or territory where the cores of long haul networks and head ends of Metro networks customers access converge. In other words, we are looking to position ourselves at the location of the greatest density of customers.

Secondly, we control the infrastructure by which our customers access us — namely the fiber vaults at the building's entrance and exit, the vertical and horizontal conduit within the building handed the distribution frame connecting the fiber to the Core Director. Third we must have a minimum of 2000 square feet and conditioned AC and DC power. Fourth and finally we use a software based inventory and provisioning platform.

***COOK Report***: Tell me more.

**Muse**: In March 2000 we filed about 50 patents on technology that we had built relative to interfaces and modeling that we did to help carriers utilize this kind of exchange. The software makes it easy for them to understand what interfaces are available at any one point in time. We also wanted to give them the capability to put their fingers really into the workings of the Core Director program where it is our intention that they will eventually be able to provision their own circuits. We have a java run web-based interface relative to the inventory today. For circuit provisioning, they will get their software hooks via open API interfaces.

All the carriers have Metasol or other provisioning platforms. We gave give them software hooks so that they will be able to use their own provisioning platforms for lighting circuits within the core Director. They are called OSS platforms. Using these the carriers can self provision infrastructure within their own networks.

What the carriers do not have is good visibility into their own interconnections with other providers. When combined with their provisioning programs, our software will allow them to see their interconnections. As a result using software of their own they will be able to see what interfaces exist with other providers. At this point they can instruct their software to interconnect with an interface coming from another provider. When they do so LayerOne watches the event, authorizes it, and then builds the interconnection out to the desired to provider.

***COOK Report***: How do I find out how much a given interconnection will cost?

**Muse**: Through discussions with the other carrier you have already determined the price. Price negotiations must always come in advance of actual circuit provisioning. It's possible that this had bandwidth provisioning Interface could also be used to bandwidth traders. People sometimes look at us as a bandwidth pooling or trading point. In reality we're not. The reason that we're not is that we don't do anything associated with bandwidth trading.

In the trading business model you set up a multi-lateral contract so that everyone has the same terms to deal with. They do

an anonymous listing which means listing the routes available for purchase. They set prices for those routes and then bring buyers and sellers together. This is the business model of Band-X, Rate Exchange, and Arbanet.

Now our platform hook could give API hooks to bandwidth traders so that they can do their standard contract. They can then give the carriers all the listings of what is available for sale. When they needed to provision a sold circuit, on our platform they could actually make it happen. The biggest problems that the emerging traders have are those of liquidity and the actual ability to provision when they sell. Carriers do not have adequate connectivity to one another. That is a liquidity crisis. Now we create liquidity by bringing the carriers together and the same device with which we do this (the Core Director) may also be used by the bandwidth seller to deliver what he has sold.

***COOK Report***: and that's a very interesting. Because when I talked to an authority in this field a couple of weeks ago and asked him how things were coming, he replied that the biggest roadblock to progress was the problem of arriving at uniform contracts and that a large part of this in turn a hung on the issue of penalty for non delivery. He added that people were really frightened by this issue because they didn't have been of a sense of what it will take for them to be able to guarantee delivery.

**Muse**: If you look at a contract for telephone service today, you will find that you are not guaranteed much of anything.

***COOK Report***: But what you are doing would make it more feasible for such guarantees to begin happen?

**Muse**: Provisioning guarantees would certainly be more attainable. If you sell something, you'd have to be able to a provision that within a defined period of time after the sale or you pay liquidated damages. We solve that part of the problem. but we do not solve the rest of the problem which is than that quality of service conditions will be attached to most contracts.

***COOK Report***: In other words while you plug the two networks together, you have no ability to do anything about how the remainder of the seller's network is run.

**Muse**: Quite true. A couple months ago the carriers walked out of one of the bandwidth trading meetings. They said "all that you want to do is to comoditize what we do — making it cheaper and more interchangeable. Why should we cooperate with you? The bottom line is that everyone is nervous as they can be about this and that, for this and other reasons, they don't have a contract yet. So we're not in the position of helping bandwidth traders trade. But, if at the end of the day, this begins to happen, it will be great solution for both the carriers and the bandwidth traders. And if it never happens the carriers will still need to get interconnected.

***COOK Report***: How you define your competition? Have you mapped out a totally unique business model?

**Muse**: We compete in four different categories. Interconnects; cross connects; peering and transit; and provisioning and trading. Looking at the interconnects, you will find that in some respects we compete with MFS, TCG, and XO but that in most respects these companies are our best customers. To some extent you can say that we compete with perhaps 5% of the local access market. This 5% is the carrier-to-carrier local access market.

The reason that MFS, TCG, and XO stay happy is that the biggest product they sell is assisting the long haul carriers to reach local markets. Although a long haul carrier could connect through us directly to Global Metro and by pass MFS, TCG, and XO, most do not. The business model of MFS, TCG, and XO focuses on reaching beyond our exchange and serving the metro area — something that we do not do. Furthermore in connecting at level one they get advantages in cost speed and reliability that they would not get with interconnecting routers at level 3.

Let's consider cross connects. Wherever you have carriers meeting you have cross connects. The problem is that these are generally single physical connections used to provision a single circuit from provider A to provider B. In a co-lo facility after a while, they tend to form a confusing tangle of wires.

And we consider our cross connects to be of a quite different nature. They do not take away any of their revenue. The revenue that we can get from cross connects is relative to the teledensity created by our shifting of the interconnection paradigm through very very cost effective optical multiplexing of resources.

The answer is that if you have a hundred carriers at a single location that is a fundamentally different situation than the one you have when you have 50 co location customers in 50 different locations within a single co-lo facility and all of them need to get to XO. In this case you are talking 50 physical connections that you have to build and test out. If you have 50 carriers connecting to each other optically, the number of interconnection possibilities is several orders of magnitude higher.

## Peering and Transit

Let's move on to consider peering and transit. While we will not facilitate multilateral peering by running a shared a media that gets in the middle of the connections, we can still facilitate peering connections the carriers have built. Of these the most critical are of course the high-speed interconnects of the Tier One providers. If you have your peering relationships already well-developed and all you need is easy access to carry them out, we are a great solution for that objective.

We have developed a particular solution for establishing private peering exchanges between Yahoo and the seven largest Internet backbones. While it is a bit technical I would describe it as a combination of our DWDM solution plus our tributary model that really gives them great flexibility. Nevertheless at the end of the day, we are really not competitive with the NAPs. People who do peering

with us simply do not do peering business at the NAPs and haven't for a long time. However the over riding factor is that we help our customers implement only bi-lateral peering. Two networks inter connect and peer and while they would generally use high bandwidth, the determining factor is the bi-lateral nature of the connection not the bandwidth. We do not facilitate multi-lateral peering between 3 or more networks plugged into a shared environment like a FDDI ring. This gets into layer 3 issues - an area that we do not want to play in.

Finally lets look at provisioning in trading. Light Wave and Enron have built pooling points to help facilitate the trades that they hope to do. These efforts however are primarily associated with city pair trading between long haul providers.

***COOK report***: What would a pooling point look like?

**Muse**: Enron has a pooling point that they have placed in our Miami Florida facility. What they do is get carriers who are plugged into our system to go through it and into their system. And then they make the anonymous transactions happen on their own frame. Enron is our customer because they could not get carriers to come directly to them. What they decided to do instead was to come to us where the carriers already are. Today Enron it does not have any people going through that platform but they hope to have people doing so someday.

In the meantime what we've done with Enron, Rate Exchange, and Arbanet is to say "why deploy your own platform when you can come to us instead and do it on top of ours?" At present we are hopeful that Enron will do just this. If they do it with us, they can have everything ready to go and not have to pay for it until they actually use it.

Now at the end of this year the beginning of next when we roll out the shared BGP DWDM ring, our bilateral peering environment will change. That ring could be used to negotiate multilateral agreements at layer one. But even though we could do this I don't think we will.

Let's consider, what can happen if Bill St Arnaud gets his way and OBGPbecomes a standard. Then customer "x" could use OBGP to directly connect its own lambda with that of customer "y". Furthermore, those customers could themselves route the results. There are development efforts underway to create really intelligent edge devices that will be able to do this routing from the edge rather than the center.

***COOK Report***: He also has a team at Carleton University working on putting code in an inexpensive switched that would sit in front of the router should the OBGP standards process bog down.

**Muse**: The limitation of a process is it that everyone who wanted to play would then have to use the same switching equipment. Furthermore I think it's likely that everyone would have to connect to everyone else in a one-off mesh type fashion. Unless that this they come to optical exchanges like layer one to do their interconnects.

In anticipation of this we will be creating an OBGP Based open lambda. Therefore he will be able to take the header information from an optical packet from one of his labeling routers on the edge. The header data will say I need to go to university X. therefore send these packets via that lambda to university X. And then other packets might be coming down the stream saying I need to go to university Y. Our device would basically just accept the information coming in from the edge and send it on. The idea is that we can take instructions from intelligent edge devices and just send them on. The problem with routers is that they have to have a memory that contains routing tables. if you can do it the opposite way and take information from the edge that says as long as you know where you need to go, you don't necessarily need to know where everything is. You really need to know only where your destination is located. Your device at the center can tell me how to get there. And I will see that my device puts you on the track on which you need to go.

***COOK Report***: Well now I understand a lot better what you're doing. I cannot believe that while you may be in 8 or 10 cities by years end there would be a market for you in a great many more. Is there really no one else who is trying to do what you doing?

**Muse**: I don't know anyone directly who is doing it. We have been out in the private equity markets raising additional capital and the folks to whom we're talking have not run across anyone else who is doing what we are. There are people who were on the periphery of what is happening but there is really no one who is doing exactly what we are today. We have 15 sites nailed down and our goal is some 30 cities total. This is a number that we think the best taps out the Telecom density that can be advantageously interconnected by our technology.

## Standardizing the Process for the Next Level and Foreign Expansion

However we think there is a macro market for the next 100 cities. Therefore what we've done this to say let's participate in standards committees and open this up as standard. We have adopted a buyer pays philosophy for our inter carrier connections. We've done a lot of other things. What we have decided is to open it all up and tell everyone how we do the things that we do, so that, if you are a co location operator in Kissimee Florida and you have a group of carriers already on your site, say 5, 6 or seven, you can help them do business with one another more effectively than they can do business today. Although there will likely be limits to the size of such a business, you will nevertheless be able to add value to your local market.

We have four sides fully open now where carriers are doing transactions with each other. We have another two where carriers are installing fiber. We are closing another round of financing and expect to be able as a result to complete all 15 sites

in about the next four months. We will turn our attention to the final 15 as soon as the first 15 are running and stable.

***COOK Report***: What about Europe and Asia? Somebody else must be doing this somewhere?

**Muse**: It is possible. But we think it is most likely to happen as a result of a confluence of events that began to take shape in United States only around 1996. Three things came together — the coincidence of lots of different networks, a lot of applications and the need for a lot of bandwidth. These three issues have combined to create a terribly acute problem. This problem is why your friend Avi has his customers calling him and saying that it takes forever to get anything because the truth is that carrier connectivity is so badly fouled up it's almost impossible to get anywhere .

***COOK Report***: Probably the situation in Europe and Asia is that the three factors are not as acute there and the additional fact that they do not have nearly as many carriers as you find in North America?

**Muse**: Absolutely. In Paris for example you don't have 150 carriers. I believe the figure would be somewhere under 20. When you get the operational factors that you have here becoming also fully mature in Europe, you certainly will need our solution in Europe as well. We do have ideas about European expansion and have eight facilities under option there among them: Paris, Marseilles, Frankfurt, Rome, Geneva, Stockholm. Last June we bought a company called Co-lo Ho which has secured us leases is in a lot of locations. For a while we did have some grandiose ideas that we would go into Europe before the need was fully mature. But, given the state of the capital markets in March of 2001, it seems likely that those plans will have to be postponed. We have made a decision to focus on the United States first.

## Afterword from a Subscriber

# What It Is That LayerOne Has Done

**Editor's Note**: Frank Coluccio is a *COOK Report* subscriber with whom we discuss technical issues from time to time. During the preparation of the LayerOne article he expressed strong interest in the subject. We asked him to react to a draft. He (fcoluccio@dticonsulting.com) is president of DTI Consulting Inc., 133 Beekman Street, South Street Seaport, New York 10038, (212.587.8150).

Gordon: allow me to get some housekeeping issues out of the way. I'll refer to Layer One, the company, using Upper Case. And I'll refer to layer one, the physical layer of the OSI-Reference Model, in lower case. Also, as a disclaimer I should further state, for the record, that I have no financial ties or interests to, or with, any of the players mentioned in this writing.

As you know, some of the reservations I initially voiced to you about Layer One's claims were either premature or without substantial merit for the optical platform they've deployed. Some of my concerns were based on dated information, as well as from my not being aware of some new ground that Layer One had broken in conjunction with their chosen optical switching vendor.

And after re-reading the interview for a third time (yes, I said third), and after doing some more diligence on the switch, I've done a 180 on some of my earlier concerns, and amended some others. But all in all I'm convinced that they (Layer One) are in a lot better shape than I had previously thought. A lot better, indeed.

If we accept what Alex Muse states on face value on this matter, and I don't know of any reason why we should not, then I'm not only pleased, but also immensely impressed with what they have been able to achieve. And I say this as someone who has contemplated - to put it very mildly - on more than one occasion the same type of neutral aggregation and switching architectures for local providers and IXCs, alike, during some earlier days.

## We Have Needed These Achievements Since the late 1980s

Our earlier efforts were at a time when optical technologies were just not up to the task, yet. In lieu of the real thing I recall drawing a future black box on one diagram, and labeled it the "Phase II upgrade," behind rows and rows of add-drop muxes for the discreet carrier feeds on the same diagram that represented the Phase I rollout. That black box in the diagram was intended to do many of the same things as the Core Director and other switching elements do, today.

In this sense I actually see what Layer One has done as a form of validation for a number of business plans that I, along with a number of my partners, had a hand in developing over the years, even if many of those never went to fruition.

One such model was half-seriously dubbed 'MECCA,' which was an acronym that stood for Multiple Entry Common Carrier Architecture, a multi-ring-based SONET platform that consisted of eight different carriers each supporting their own nodes, and each supporting their own cross-connects. The primary challenges were designing the fiber optic meet point and SONET interoperability, which apparently LayerOne has not had any trouble with. Handoffs were to be accommodated by add-drop mappings. We suggested this design to one of the Stock Exchange's carrier subsidiaries in 1993.

We perceived this to be the next step up from the simpler DACS model that was used in the Securities Industry Association's community of interest networks of the day, such as the SIA Data Network, or SIDN. Today, I would suspect that such a design might revert back to a form

of resilient packet ring architecture, if not one using a 10Gb Ethernet approach. We do work in cycles. LayerOne talks in terms of effectuating a paradigm change in the way carriers interconnect. Certainly it is bringing about a vast improvement. But in doing so it is building on momentum that has been there for more than a decade.

From the late eighties to the mid nineties the digital cross connect system (DCS) was to play a heavy role in each of these models, since it was at the time constantly being upgraded to accommodate increasing line rates, and eventually accommodated native SONET rates, as well. Consequently, since the DCS accommodated the highest line rates and subrate densities of any platform, it was also the most cost effective option available for the purposes at hand.

In principle, the DCS did then the same thing as the Core Director now does; only it did so under a less responsive mapping regimen, and it took about an eon longer to accomplish, in relative terms. But the DCS did, and does, just like the Core Director, radically reduce the number of lines that were required to support a very large number of connections. It accomplished this by eliminating a great number of lower speed voice grade circuit provisions, including modem lines and stat muxes, replacing them with T1 and T3 lines and digital data services.

One client went from over 20,000 analog copper lines to just a few diversely routed fiber loops (plus some holdout copper lines for wet circuit battery in the event of a total fiber outage), in about a year. Paradigm shift? I don't think so. It was instead a dramatic improvement in logistic- and administrative- handling, with a commensurate decrease in circuit costs. That was radical in its time, but it's altogether commonplace today.

## LayerOne Offers the Industry a Reproducible Template

I applaud LayerOne for executing their plan at this point in time as they have, and in so doing, apparently, gaining a first-to-market advantage in this space. From what I can see they've put together a reproducible template that exemplifies what I see as the next significant step up in the local exchange model, even if they don't see it quite as such, yet - the extent of which has been made possible only recently through advances in opto-electronics and the extremely dense integration of very high speed switching applications in silicon.

I would like to add that other fiber optic exchanges of considerable size and density exist in such venues as IXs, peering points, NAPs, colos and hotels, as well. Most notably, and aside from the more common "meet-me" kinds of rooms that are typically found in the just-mentioned exchanges, there is the agreement between Verizon and Metromedia Fiber Network, for example, that allows MFN to aggregate CLECs over MFN fiber in Verizon serving areas on an optical colocation basis on Verizon's own frames. To the best of my knowledge, these are still passive cross-connects at this time, although it stands to reason that they would upgrade them at some point to o-e-o in a way similar to what LayerOne has done, or simply using their metro WDM ring capability as they do for enterprise customers, by placing CLECs and ISPs on their own lambda allocations.

I think that it's also very telling to note that LayerOne claims these two carriers (Verizon and MFN) as customers of their own. Especially Verizon, who I would have suspected to be among the staunchest holdouts of all, but they're in there, too.

## Some Key Issues

Allow me to summarize the issues that stand out in my mind from your interview with Layer One's Alex Muse:

First, Layer One claims that their Optical Exchange represents a new paradigm. I happen to disagree with this characterization, primarily because they are actually leveraging all of the time-honored capabilities of time-, frequency- and space-division multiplexing (SONET/etc and wdm, respectively), only they are doing this in a far more efficient and dense manner, to put it mildly.

In the strictest sense of the term, therefore, Layer One does not constitute or foster a new paradigm, per se. Instead, it has capitalized on recent technological improvements and presented a much-improved 'paradigm' than the one that existed before it. If the manner in which they do this can be interpreted as a paradigm shift, then so be it.

Secondly, the exchange hardware from Ciena "talks to," or inter-operates at the physical layer, with other devices from at least three or four other optical vendors. From what I have read, there needs to be some additional level of assurance from the vendors involved that they will guarantee that those boxes will in fact always talk with one another. At the present time, and I'm not certain about this, but I'm inferring that no such guarantees exist. I've also inferred that the capabilities that Layer One demonstrated were possible, were not capabilities that their vendor, Ciena, would have initially guaranteed, much less expect LayerOne to attempt to accomplish.

I would expect that before they placed the reliability of their customers' (other carrier's) links in the balance, that they would first want to get such assurances covered in their procurement and service contracts. In fact, according to LayerOne, Ciena was surprised to learn that their own box was actually able to talk with competing optical network elements. I would sleep a lot more comfortably at night if I were the guy sitting in the LayerOne NOC if I had contracts with sharp teeth in them with the vendors in question, stating that their optical network elements will talk to one another.

I'm not questioning LayerOne's reporting on the fact that they achieved interoperability. On the contrary, I have every reason to want to believe that they have, especially as it relates to the MAN venue. I think that WAN optics, such as schemes now being used in ultra-long hauls, will present another set of issues at some point, but that's another topic. I am, how-

ever, suggesting that they need assurances from not only "their" vendors, but also the other vendors that their customers use, as well, stating that they will be able to talk to one another in a consistently compliant manner across a wide range of operating conditions. I would strongly suspect on the basis of other measures they've apparently taken that they've addressed these issues, as well.

## Though Situated at LayerOne Core Director Also Carries out Higher Layer Functions

Third, while LayerOne states that they are operating primarily at the physical layer - indeed, their name itself causes one to conjure up images of nothing BUT the physical layer - one must look beyond the words that make up the name, and peek behind the silicon curtain to examine what is happening in the switching fabric and on their backplane, in order to gain a better understanding of how they do the things they do.

True, their I/Os are layer one, as are every other carrier's at the port level in every SP, central office location and exchange point. If we believe in the OSI-RM, there is no other way to enter the physical port of a box than to go in at layer one. However, there's a bit more than glass works going on inside.

Ciena and others will do the same things that have been done inside digital cross connects (see clarifications below) with regards to SONET/SDH, add-drop, grooming, etc., until lightspeed switching at the packet level is both possible and economically viable, or until end users and carriers, alike, cease to use those formats to define their units of currency. Namely, OC-3, OC-48 and OC-192.

Here are the clarifications I alluded to earlier: The main difference is that the newer, primarily-optical boxes are driving these processes deeper into their own backend switching fabrics, as opposed to using discreet boxes, such as external back-to-back add-drop muxes, to achieve the same end. By doing this, they hide all of those ugly and expensive (to buy, to house and to administer) boxes from view. The Core Director, as a case in point, along with a growing number of other vendors' platforms, does some of this now and will be doing more of the same, in addition to the more-obvious lambda-centric applications, such as MPLS and MPLamdaS switching and routing schemes, too.

Those are the constructs that you need to focus on in order to fully appreciate how their platform takes an "optical" signal at the OC-192 rate and then splits it up into a so many lower speed "optical" tributaries, with each of them mapped to their own "optical" lambda port. Behind the curtain it's all taking place electronically, and at various sublayers within the stack. In other words, one has to un-do all of the optical related imagery that we use for mental props when trying to figure these things out for the first time, and do a simple walk through of the more mundane steps that take place in the electronics within the switching fabric. Hint: just sit back and visualize all of the back-to-back legacy boxes that are being displaced.

Fourth, and this is in some ways directly related to three, above, future offerings will include lambda switching (a la MPLambdaS), and other forms of layer 2 processes. At one point it's even mentioned that Optical Border Gateway Protocol, or OBGP, will, or might, be supported. These claims would appear to fly in the face of what was stated to be a purely Layer One paradigm. Or, does it? After all, every Telco, CAP/CLEC, CLEC, ICP, ISP, NSP etc., must interface at the physical layer, or layer one, of the OSI-RM. There is no other way to get there than by using layer one, no matter what rides atop.

## Paradigm Change?

Getting back to our earlier question: Does the LayerOne optical exchange model represent a new paradigm or "paradigm shift?" With all due deference to Alex Muse and Co., my preference would be to avoid such verbiage. Most of what is called a paradigm shift these days can actually be traced to earlier iterations of the same model of architecture, which were based on the same fundamental principles that the new version uses.

So, if we had to use the "p" word at all, I'd call it a paradigm enhancement or improvement. Although, their horizontal business model, which I could go on about at length for its brilliance is a radically effective one, which capitalizes on new capabilities at just the right moment in time to effect superb market positioning. And all the while capturing and harnessing many of the ideals that have been sought after for the past hundred and fifty years of networking, beginning with the telegraph. Has this ever happened before? Certainly it has, many times, but in the past it's always been to a lesser degree, the same as could be stated about each previous improvement. And that's my point. On a global scale issues of the existence of at least of some degree of backward compatibility will continue to be a concern for an new technology that expects to make an impact. Here the issue of compatibility is the bandwidth unit of currency (OC3, OC48, DS 3 and so on.)

The model keeps improving with time, by degree, but not by anything radically new in principle. When "radically new in principle" happens - and I'm not referring to something as bland as sending voice over packet or anything like that, which I've also heard called a paradigm shift in the past - then we'll have the potential to disrupt all that came before it. And very few want, or can afford that, right now. For that matter, as complexity increases, it will become even more difficult to do in the future.

## A Note to Subscribers

Effective with this issue we have switched from Page Maker to Quark Express. We are using auto leading which we hope you will find makes for increased ease of reading.

# Changing Bandwidth Provisioning Models in Metro Area Fiber Markets

## Net Access Chooses Acquisition of Dark Fiber and Self Provisioned Circuits Over Purchase of Shared Gig Ethernet Prices of Telseon and Yipes! Viewed as Too High and Cogent's Business Model Seen to Be Unsustainable

**Editor's Note**: While still in college Avi Freedman founded Net Access, Philadelphia's first ISP in 1992. In the COOK Report's basement in 1995, he put in a 56k-connected "POP." However, at the same time, he discovered DS3s, backbones, and routing. Consequently, the POP never got beyond the 100-pair copper cable from the pole to our roof. Since 1995 Avi has been one of the most widely-known advocates for small and mid size ISPs. In 1998 he joined AboveNet as VP of Engineering and, in late 1999, he joined Akamai as Vice President and Chief Network Architect. As he put it in February 2001, his weekday job is still with Akamai and his weekend job is CTO of Net Access of which he is still a principal owner. We interviewed Avi on February 8, 2001. He speaks in this interview in his capacity as Net Access CTO. He last appeared in our June 2000 issue in a long interview on Akamai.

***COOK Report***: By way of background please tell me about the current status of Net Access and its infrastructure. I gather you still own the network even though you are not running it on a day-by-day basis and that you have infrastructure in a number of East Coast cities.

### Net Access: Infrastructure and Business Model

**Freedman**: For a couple of years now Net Access has had a T3 and OC3 network that goes from Boston to New York to Philadelphia to Baltimore and to Washington D.C. It is part of the network that has transit from AboveNet and Level 3 and now OC12 transit from AT&T and Sprint.

***COOK Report***: What is the business model of Net Access?

**Freedman**: It's really a three part model. One is IPconnectivity. There are 80 ISPs and Internet centric companies that are wholesale customers of Net access. There is an enterprise market, which at this point, is some number of "n" thousand of enterprise, connectivity, and dedicated co-location customers. Finally, there are a number of content providers as well as about 5000 dial-up users primarily in Philadelphia. Although we began with a dial-up network, dial-up is presently a very small part of Net access revenue. Net Access is now $10 million a year company and growing. We offer three core services. One is leased lines. The second is co-location. We have opened the first 15,000 square feet of a co-location Center in the Conshohocken area of Philadelphia - one that is costing $6 per square foot per year and is enabled by dark fiber. We took possession in May, went live in August and had the "grand opening" in October. The center, called EarthStation Netaxs, is in Lee Park, in Conshohocken, PA.

***COOK Report***: I'm beginning to understand how you could be in a position to be a customer for some serious infrastructure.

**Freedman**: We are not yet in the MFN in top 25 but we hope to get there. Now our third service which we had not really thought would be big but is turning out to be, is selling circuits. The model is simply if you want a circuit (OC or DS3), tell us and we will turn it up for you from any co-lo that we are located in to any other co-lo that we are in. A lot of people talk about buying lambdas. But you actually don't want to fill lambdas because they can seriously do you harm if their frequencies are off. And generally we would not sell you an OC48 lambda per se. What we would do is sell an OC48 that just happens to take up a lambda.

***COOK Report***: Where is Gigabit Ethernet, in the overall scope of what you're doing?

**Freedman**: Gigabit Ethernet is something that we've actually found that people don't want primarily because of companies like Telseon and Yipes!. The fact is if I give you an OC12, you know you are getting an OC12 and there's nothing I can do to short change you on the matter of how many bits you can pump through it. You know that you have your full circuit even if it is being time sliced and sent over my WDM infrastructure.

But if I give you a Gig E, you have no clue whether that's a VLAN on a big trunk or whether it is a real Gig E dedicated all to me, because the technology allows you to sell 10 people a GigE, put a 10 meg Ethernet in the middle and then sell you all Gig Es at the other end. Furthermore there is no 10 gigabit Ethernet switch today, so if I tried to do aggregation, you would be unable to tell what I had done. You might think there was a full GigE there, but you really couldn't prove it.

***COOK Report***: The point you were making is that there is a very big difference in having a gigabit Ethernet and your own lambda for in the latter case you know you have the full 1000 megabits. But if you have a gigabit Ethernet on a shared transport mechanism, then you have no

way of measuring exactly what you are being given.

**Freedman**: There is a great analogy between the people who are doing in the managed gigabit Ethernet " stuff " and ATM and frame relay garbage that was forced upon ISPs some years ago. Basically, unless you can see the infrastructure, you have no way of knowing what is in the middle.

***COOK Report***: What does it take to see the infrastructure?

**Freedman**: It takes the willingness of the provider to put up graphs showing the link utilization of all the points on their network starting with the end at which you connect and ending with the other end at which you connect. The tools MRTG or RRDtool (by the same author) are easy to come by. The tools show over time what traffic is going from device x to device y. The device could be either a switch or a router. So if someone says he will give you 100 megabits and he will put you on with nine other 100 megabit customers on a Gig E trunk, then you can go and see that indeed this is what is happening. Or, if it is not happening, you will find that out as well.

***COOK Report***: Doesn't it come down to the point that that if you own your own fiber from point A to Point B, you can put a gigabit Ethernet switch on each end and know for sure that you have a Gigabit worth of bandwidth? However if you do not own the fiber and somebody else it supplying the GigE, there becomes a problem of trust. Is that essentially where we are?

**Freedman**: Yes. And let me clarify one issue: that is what is meant when we say "own." By "own" I mean have the right to provision fiber that there is nothing else on. In other words, while I don't own all the fiber I get from MFN, for 20 years no one else may put anything on it except me.

***COOK Report***: I think that you are talking about an IRU?

## The Dark Fiber Business Model

**Freedman**: There is a significant difference between a 20-year monthly lease and an IRU. The IRU is a capitalized instrument which you depreciate. The 20-year lease is just like a monthly telco expense. With the IRU they typically give you a discount over what you would pay for the total period. You pay for it upfront or finance it. And the carrier gets the money in a lump-sum up front. And you get to treat it as a capital asset on your books rather than as a liability. If I'd buy one of your 20 year contracts to get fiber, for me it is booked as a liability. If I buy an IRU, it becomes an asset that I depreciate. At the end it is all the same. I have the fiber for 20 years whether I pay for a monthly or whether it is an IRU. With a lease, in every case, no matter whether it is 3, 5 or 20 years, I pay a monthly fee.

***COOK Report***: So you could actually commit to a monthly fee for five to 20 years because you know that while the price of bandwidth will change through DWDM, you will be able to put more and more bandwidth on that strand of fiber.

**Freedman**: I can certainly think five to seven years on fiber because indeed I can add more bandwidth. The cost of generating a lambda (on which you can put SONET or GigE) from a single strand of fiber is certainly headed downward — while the direction of the cost of IP over fiber is less certain.

***COOK Report***: Have you just done 3, 5, and 7 year contracts?

**Freedman**: No. I've been involved in 20 year leases because that seems to make sense. For example let's say you want to be Mr. Network in New York City. But at some point, New York becomes ruled by the Green Party which decides there will be no more digging in the streets. What happens then? How do you get the fiber you need? Consequently I would rather have the certainty of knowing that I would have the fiber. Now the fiber itself may not last 20 years. But then you will have to replace it and that will be their problem.

***COOK Report***: People seem to think it will be good for least 15 though?

**Freedman**: Unclear. Given the power levels of some WDM gear, some people think it that the equipment used with the fiber could be damaging it. The big difference is the gear that's used with the fiber, not quality of the fiber itself. The issue is that if you drive your laser equipment at sufficiently high power, could it actually damage the fiber itself? There are people who are going to start doing that relatively soon.

***COOK Report***: Why would someone do this? What would be the motivation?

**Freedman**: The difference between more expensive long haul equipment and cheaper short haul equipment is pretty much the quality of the lasers. If your lasers are super high-quality you can go longer distance with less power. But if you're trying to do something strange like Net Access which is using metro fiber gear to amplify your bits over a long haul, then you can get into the danger zone if you push it. (Now we're not pushing it but we could.) Everything is engineered to be roughly 28 DB between amplification huts.

We are within 28 db on the equipment we are using. But there is new equipment from ONI and Zaffire that is promising four to 600 km range with cheap lasers. Which many people who know more about the fiber physics than I do - and as you can imagine spending a lot of time near MIT I meet a lot of people who know physics - claim is potentially damaging to your fiber and could lower its lifetime toward the five-year range. This is especially true for long haul fiber but, but to the extent that metro fiber is part of it, it simply is part of the same over all network.

## Pricing: OC Pipes Versus Dark Fiber

Now Net Access has a room full of Cisco equipment to power the fiber network that we have purchased from an MFN. (The room is because it's being set up in a lab and not yet deployed.) The network will be run from Conshohocken, but the Cisco gear will be and is stationed throughout the physical extent of the network itself. There is a ring that goes around the Philadelphia area. We have lit

one segment of it using $5,000 Cisco gigabit Ethernet switches and attaching generically named ZX GBIC laser electrical interface modules to the switches. The GBIC is a standard and the card that Cisco sells is one made by someone else which Cisco marks up and resells. So for $8000 on each end you can do 70 kilometers of gigabit Ethernet.

But then it gets trickier. If you look at how to light fiber, there is this equipment which only goes 70 kilometers. But then there is coarse WDM made by people like N-BASE Xyplex where you can put four or eight lambdas on a fiber pair. Finally there's the normal DWDM gear which is stuff that can do 32 to 256 or more lambdas. There are two varieties of this. The metro gear and the long haul gear. The metro gear can go typically no more than 28 db or roughly 100 kilometers without amplification. These boxes all cost somewhere between 20 and $40,000 per lambda per end.

***COOK Report***: Suppose you needed some 20 to 35 lambdas from point A to Point B. Is it more cost-effective to invest in a single pair of fibers and more expensive DWDM gear or to take the coarse WDM gear and use several more strands of fiber?

**Freedman**: It is often more cost-effective to buy more fiber than to pay the large depreciation on the DWDM gear. Based on customer demand, because Net Access needs to be able to sell circuits, it doesn't make sense to burn a whole lambda for one circuit and do coarse WDM there. In other words I can put 16 OC3s on an OC48 lambda.

Our first thought was that we would extend our fiber into the co-location facilities and sell people IP connectivity. But we found that our customers were becoming apoplectic about their relationships with the phone companies which were taking six months or longer to provision circuits. Still thinking IPwe asked: to their Internet connection? And they said no. Between 111 Eighth Avenue and 8100 Boone Boulevard (MAE East). In Manhattan 111 Eighth Avenue is where people are now because at 60 Hudson is both full and expensive. Of course 111 Eighth is almost full now as well. But while there are a lot of people there, the cost of room to room cross connects is pretty reasonable.

There are a lot of international providers who are in New York and want to connect to down to D.C. When they find out where we will take them, they ask how much we will charge them for a cross connect to some other building in the same city? So we figure the cost of the gear and profit and tell them $3,000 a month. And they reply: that is only half what the local exchange carrier would charge. They say, when can you do it? We reply that we can do it within a week or two after the fiber is in. Then they say do you mean that we can continue to order circuits from you? Sure we reply. Once we are up our provisioning time should be no more than a week.

Now there are people like Enron who spend a lot of time and a lot of money on software that will do circuit provisioning in a minute. I think that there is clearly a happy balance out there somewhere between "it will take me a week" and "it will take me a minute". The number of people who need a circuit in a minute and would know what to do with such circuit if they had it is relatively low. Of course it may be that you don't want to order in OC12 for only 15 minutes. But on the other hand you can order the local loop to someone and get temporary bandwidth for 30 days.

***COOK Report***: When they say they need a circuit between New York and Washington, who is in a position to sell it?

## Provisioning New Circuits

**Freedman**: Qwest, Level 3, Sprint, WorldCom, Broadwing, Winstar, XO. Those are all potential candidates.

***COOK Report***: How long does it take?

**Freedman**: 3 to 6 months. At Akamai, from one of the carriers whom I just mentioned above, we ordered two OC12s and four OC3s about seven months ago. None of them have been installed yet.

***COOK Report***: Why are they so slow?

**Freedman**: I simply don't know how it can take them so long. Other than, as a friend of mine said: "they have difficulty with fecal aggregation" - in other words with getting their s___ together. Net Access used to be able to order a T-1 and get it installed in a day and a half in the Philadelphia area in 1995. It currently takes months to get a T-1 in from Bell Atlantic (now Verizon). Verizon wound up using Nynex's methods because they thought that the Nynex people who did leased lines were more wonderful than the people who did leased lines for Bell Atlantic and that of course was simply not true.

Now WorldCom (MFS) used to be able to give you a FOC date - Firm Order Confirmation. Now it is known as a "FOC you" date because it is so unreliable. It used to take only a day or two. You would put the order in and get the date and know that you could count on no more than 30 days from that date. But now WorldCom decided to use the WilTel circuit ordering system which has to be the worst in the industry. They then wound up taking a month to get you a date by which they might have it installed.

***COOK Report***: Is the reason for this a huge expansion in the orders that they must deal with?

**Freedman**: I don't know for sure. Regardless of volume, it should not be an insolvable problem.

As far as the core aggregated Ethernet business model goes, the fact that it doesn't make any sense might mean that their investors will get some return because other investors will buy their software. I know that there is at least one company that has gotten venture funding just to make software that will let people provision.

***COOK Report***: I think Williams makes provisioning software.

**Freedman**: Well, as to how fast Williams could turn up a pipe, I must say they haven't given us pricing that is good enough to make us want to try. This doesn't mean that they are bad company. It just means that we have not been successful in extorting good pricing from them. Seriously, when they do their pricing, a lot of these guys don't seem to realize that their potential customers can go out and get their own dark fiber if their prices are not good enough. I should say that the carriers I know find Williams dark fiber pricing attractive, though, and besides normal building termiantion delays, they seem to be able to deliver reasonably well. But once Williams delivers dark fiber, their job is over in terms of the circuits that ride on that dark fiber.

## Economics of OC12 Versus Dark Fiber

Let me put it this way. Net Access has tried to order in the past perhaps 8 OC12s and from its experience has determined that lighting your own fiber is cheaper if you are already acquiring such. If I wanted OC12 in one of my cities, most of the companies had been talking about charging me somewhere between six and $14,000 a month depending on the distance between its start point and termination point. They all use the Bell model, which means there's a termination charge per end. If it goes outside the CO there's a separate charge. Finally they charge you some rate per mile per month. This rate generally depends on the length of your contract. It is basically the same way that T1s and T3s have always been priced. By the way, if you don't take a five-year term on that OC12, the circuit becomes about $25,000 a month.

***COOK Report***: So you do the math and find out that, if you already have some of your own fiber, you might as well use that or get more?

**Freedman**: For about that price you will find that you can call anyone who sells dark fiber (MFN, Level 3, and others) and you will find that for about the price of one five year contracted OC12, they will sell you three year point-to-point fiber pair allowing you to connect what the OC12 would have. Furthermore, if you want to double your investment you will find that normally you can extend your fiber pair into a ring around a central downtown area. For example, in a city with a fairly concentrated downtown area you will not pay more than $20,000 a month for a fiber ring around that area. Now if you were talking about Boston or San Francisco it would be more than the $20,000 per month. But it wouldn't be double even though there's a lot more distance involved.

So let's say that for $6000 a month or $8000 a month I'm going to get my dark fiber pair from Level 3. While I like their pricing, it doesn't mean I trust them when they assure me that the ring they've sold me is redundant. They are after all a telco-albeit disruptive telco whom everyone else in the telco world hates. Therefore, I make sure I don't use them for everything. I do have OC12 s from other people than via dark fiber carriers at both Net Access and Akamai. I should add that Level 3 here is an example, as dark fiber pricing seems to be fairly similar all around in the metro area.

***COOK Report***: Are you responsible for provisioning Akamai's network?

**Freedman**: Everything that touches Akamai's network is my responsibility. But Akamai and Net Access don't compete so there is no problem in my wearing both hats.

And the network is huge. Akamai currently handles between five and 10 billion transactions per day. In compressed form these transactions take it 360 GB and that we must send back and have to deal with in real time. This is how we believe, bigger than almost any telco billing system. And this is complex but it's only about one-tenth of the total complexity of building a content distribution network in which everyone thinks only have to do is buy a couple of boxes and put them together.

***COOK Report***: So you have a need for this infrastructure. What did you find out when you went out and tried to purchase it? I gather you found out that people wanted to buy even more than you thought they did.

**Freedman**: Yes. Net Access has been selling IP connectivity and co- location related to that.

***COOK Report***: But now you are moving down on the protocol stack into a fundamental transport.

**Freedman**: We are having to do this because there is a vacuum of people selling what other people want to buy. And by people I mean smart people who actually understand how other people have gotten harmed in the past by purchasing certain services and people who also understand the economics of what their options are. What I'm talking about are the options of buying dark fiber for people who have been shortchanged by buying aggregated services blind.

***COOK Report***: And people who have no means of judging or themselves the quality of what it is that they're paying their money for so they don't know what they're going to get for the money that they do pay out.

**Freedman**: True. If you are talking about layer two packet delivery, your quality is does it get there or does it not. It is not as much of an issue as would be IP quality at layer 3 where you would want to know what your throughput was. Let's say I choose to light up a piece of dark fiber and put a gigabit Ethernet switch on either end. I could turn around and sell you a gigabit Ethernet circuit on that. If you are the only one to whom I sell such a circuit, you are happy. But what if I put a hundred megabit per second customer on there? Well now you need more than nine hundred megabits per second, you are not so happy. Now what if I say am going to go out and sell to everyone and their uncle and I sell you a gigabit Ethernet with only one megabit behind it? And I say that when you want to you can go on to consume a full gigabit.

***COOK Report***: There's a lot of multiplexing going on and you don't know how that is going to shakeout.

**Freedman**: You know there has to be multiplexing going on because no one

has a ten-gigabit per second Ethernet switch. In other words if I can do a 10 gigabit Ethernet and then sell ten one gigabit Ethernets off of that, and that is OK. But all I can currently get on a gigabit Ethernet circuit is a single gigabit. Therefore if I'm selling gigabit Ethernet and I sell such a circuit to more than one person, by definition those to whom I sell are not going to get the full gigabit. Now people could be doing strange things to take layer 2 over OC-192. But reality says probably they're not.

## Issues of Scale and the Economics of OC-192

Why? Because they don't have enough money to do this. They don't have enough money because OC-192 cards are $250,000 per port whenever router you are looking at. Furthermore this figure does not include the cost of the router. Now how many OC-192 ports can you by with a hundred million dollars? You could buy 400 OC-192 ports with the $100 million that Telseon just raised. In other words 10 ports in each of 40 metropolitan areas.

***Cook Report***: Is part of what I'm hearing that at your particular market level there is an opportunity of making a more effective use of a small amount cash when you are dealing with OC12 and nothing larger because the cost of bigger circuits is simply out of sight? While the cost of generating OC12 s and handing them out to customers, compared with OC-192s, is very economical.

**Freedman**: Yes if you can just get more fiber. If you can do 32 OC48s and get another fiber pair, do 32 more OC48s, it can be cheaper (especially if you're trying to make your capital last) to just do this. In other words the most cost-effective course is to burn the fiber pair for now and then later go back and put other gear on.

***COOK Report***: I have the impression that your economic model for Net Access from the very beginning has been one of self-financing. Have you been able to maintain that method of operation?

**Freedman**: Yes. Net Access has grown without outside investment or a huge debt. It has grown by the reinvestment of its revenue stream. As you get into optical gear, you do wind up leasing against your expected revenue stream which it is not quite the same thing as debt. To the extent that you are leasing you are covering it from your revenue or from cost savings in your network. The entire dark fiber infrastructure that Net Access is putting into place is no more expensive than the infrastructure that it is replacing. But the replacement gives us "N" times the capacity that we had when we were just buying DS3s and OC3s. And now we can go sell circuits on them if people want that in addition.

***Cook Report***: You've been working on this for how long?

**Freedman**: About six months. It took awhile to get some of the buildings lit. Some of the edge of buildings still aren't lit because of permit issues. But the core buildings are lit.

## Leveraging Infrastructure into Quantity of Strands

***COOK Report***: Was it a year ago then that it began to become clear to you that the best course of your expansion for the long run would be to pursue this model of acquiring fiber?

**Freedman**: Oh absolutely. Dave Rand who was my CTO at AboveNet first pointed out what I think is a deep truth. QoS should mean "quantity of service", not "quality of service". Or as recently amended, QoS should stand for "quantity of strands".

***COOK Report***: In other words you belong to the "throw bandwidth at it" side of the debate.

**Freedman**: Yes. Or to the school that says QoS the way most people mean it is that "if you pay more money to me I will drop the other guy's packets first". Now this brings us to the Cogent discussion and here I don't know how to describe how badly the math doesn't work on that model.

Peter Christy recently did an analysis and said "If I had 400 fibers I wanted to light in a Metropolitan area and put IP on - not just light over DWDM - that it would cost me $4 billion to put IP on both ends of those 400 fibers running say 32 OC-192s on each fiber". That's just the thing. You must realize now that IPbegins to be a real cost at those speeds.

But to get back to your question. Yes. It's been clear to me for about a year because I saw AboveNet making very good use of dark fiber for private interconnects and, for backbones and circuit replacements. It certainly was the goal of AboveNet to get entirely changed over to dark fiber. And that is why AboveNet pursued MFN and to say that it really made sense for both companies to come together. Consequently it was probably about a year ago that Net Access started asking "How much it would cost to do this"?

***COOK Report***: For about past year you been asking questions and for the past six months you been doing some implementation?

**Freedman**: Yes. I had been looking at a offerings from some of the players like Telseon and Yipes! And had immediately been saying: "Well you know, that's pretty retarded. Why would you think that I would want to buy that?" I had been talking to about eight different fiber gear vendors who all have annoyingly similar pricing although they are all configured differently.

I have the Uber Spread Sheet of Doom which actually lets you play around with how many circuits of what you want on what different topologies to show you how much each vendor's gear will cost.

***COOK Report***: Where did that come from?

**Freedman**: I built it. I can't give it to you because it uses all nondisclosure pricing. I'm going to write a Boardwatch article on fiber gear because it's all very simple. Except that all the fiber gear vendors are

very annoying. And they're even worse than router vendors. They want to tell you: "Give me the configuration and I will tell you how much it costs." And we say: "No. Tell us what parts do and how you put them together and then we can build our own configuration." The problem is sometimes even their own engineers don't understand how they put the parts together. We actually had one of the very largest of the fiber gear vendors who wanted to hire a couple people from Net Access just because we could correct their quotes.

***COOK Report***: Is part of the disconnect that you're seeing perhaps that these companies have hired MBAs who tend to work on abstract principles rather than knowledge gained from the trenches in actually running ISP operations?

**Freedman**: I can't speculate on why these companies evolved the way they did. They obviously have technical people. I think the decision may have been that five or six years ago people wanted to buy ATM or frame relay and now I want to buy a similar thing on a gigabit scale - not having realized that the pricing which they are claiming they have to get in order to be able to stay in business makes no sense.

***COOK Report***: Somewhere there must been a disconnect between their marketing people and those in change of the technology?

**Freedman**: My assumption is that the price that will turn away business is the price that they have decided that they must get in order to stay in business. It sounds like a simple concept. They could say we are simply not going to bargain because it will lead to price erosion. But if you are a new start up company, the price at which you are willing to turn away someone with cash in hand is probably some number that you figure does not give you cost plus some margin. And it is probably not 500 percent margin because these companies are about to go out of business if they don't get more money. The point is if they want to sell me something that costs more than dark fiber, maybe they just didn't analyze their market and understand that getting dark fiber would be an alternative to buying their products.

Maybe they thought that MFN would stop selling to people? But selling to enterprises and ISPs is MFN's core business. MFN sells using a return on investment calculator themselves.

***COOK Report***: And Level 3?

**Freedman**: Level 3 does sell dark fiber but more on a wholesale basis. There are other people in various markets who sell fiber as well. These include utility companies and there are some smaller fiber players. 360 networks in some places. Enron brokers dark fiber. GPU in Pennsylvania sells dark fiber. Amtrak sells dark fiber. And there is even a web site that will tell you who sells dark fiber in a particular geographic area which interests you - http://www.fiberloops.com - per fiber mile Metro fiber is more expensive because, though, because it is more expensive to install.

***COOK Report***: One has to wonder where the business models for Yipes! and Telseon came from?

**Freedman**: I don't know. I didn't talk to them about how they started. I know that every time I've been hit with a "Hey would you like to buy this", I respond that it doesn't look very interesting because we can extend our own dark fiber to those buildings. Or that the customer could just to buy another pair of fibers that has nothing to do with our network. In terms of buying I represent two different contingencies. For Net Access if I am going to provisioned a circuit to a customer - in other words if a customer wants to get a 30 megabit capacity circuit to me and Yipes! wants to charge me $3000 a month but I can extend my fiber there for $2000 a month and get $500 a month as depreciation on equipment, it definitely makes sense for me to get my own capacity there because Net Access always shows its customers link utilization. I know if I put a gigabit Ethernet in there and I give them 30 megs they will be able to see that I'm not oversubscribing it.

***COOK Report***: So you taken a lot of lessons from AboveNet?

**Freedman**: Yes. AboveNet did a lot of things right. That's why I wanted to work there. On the other hand, what I have learned from being at Akamai is that content distribution is not nearly as easy is it seems and that you should resell Akamai if you want to have a hope of having a good commercial product.

## An Issue of Scale: Hacker Ethic Meets National Corporate Roll Out

Going back into the business proposition of Yipes! and Telseon the idea that I would want to run an aggregated Ethernet for the same amount of money that I could install my own fiber and circuit makes no sense to me. And aggregated Ethernet is simply not attractive because I don't know who else is on it. When I ask them they say our best practice is industry standard . . . mumbo jumbo. . . and that we have double the capacity so that if half the network goes down, we don't oversubscribe it. And I say "OK, show me". "Our best practice is that we don't do that." And I say "How do I know that?" And they say "Because we said so. " In doing this they are operating like telco. And telcos lie because they don't even understand what customers asking sometimes. Frankly they can't know. Most in telcos did not put all their fiber in the ground so if you ask them, they can't tell you where it goes.

***COOK Report***: In other words, if you want to promise a quantity of service, you better own your own infrastructure and have control over it so that you can deliver what you have promised.

**Freedman**: Yes. And there are ways of leasing fiber and finding out, but the bottom line is most telcos don't do it. If I get across country circuit, a year ago, the chances are it would have gone through 10 different segments. Today there are a couple of vendors you can give you an OC12 across country all on their own but

not very many and not along all paths. Furthermore of the 10 paths on a circuit the chances are that at least two were lies. In other words they don't actually go where they state they do. What you need to do when you take the circuit is every once in a while take it down and tests the latency of each of the pieces and see whether they can possibly go the way that they say they do. We had a San Jose to Chicago Circuit delivered to us with a hundred and 20 millisecond round trip time which is how long it takes to go from San Jose to London and back. We made them test each segment and found out what was wrong. Now the problem is when you order a circuit, if you don't have an SLA (Service Level Agreement) on latency in your circuits, which most people don't, they can simply say we don't care how slow it is, there are no bit errors and then you are out of luck.

I have two broad reactions to what is going on. One, with Telseon and Yipes is it doesn't make sense to me. Two, with Cogent I can prove to you that they will die.

Cogent is very "strange". ISPs are buying from Cogent but it is a situation where, although they are buying from Cogent, they are not really relying on it because they secretly of course both hope and know that Cogent will die. The ISPs that are buying from them are selling what they are buying for much more than the $10 per megabi they are payingt. So that situation is very strange.

But Telseon and Yipes! are going to have to adopt a model that makes sense meaning that if you sell me a 100 meg circuit, it better cost less than half of what i could buy dark fiber for or I will just go and buy and dark fiber and have infinite capacity. Or I will charge you are reasonable amount that is greater than half the cost of dark fiber but I will also give you view into my infrastructure.

Or you can say you are going to charge me a quarter or sixth or an eighth of what dark fiber costs to give me the hundred meg circuit. Next you tell me that you have redundant capacity which I might not have and yes with my way there are more points of failure because I have all this complicated equipment and complicated software. Then you may point out that you have an extra capacity for multiple fiber vendors and pledge that you will show me the status of your infrastructure, we have a basis for a conversation.

But it becomes much more difficult if you say I'm going to charge to two-thirds or even one half of what dark fiber costs and I'm going to make promises without back up of how much capacity might to exist or might not exist between points A and point B and I'd like you buy it from me. And then I say "How much is an OC48?" and you have told me that the cost that is more than twice what dark fiber costs. I say that I don't know that that capacity is actually there, or even if it is why would I pay twice as much as dark fiber costs for one 32nd of the capacity that I could otherwise receive if I lit my own fiber?

This is especially true when I don't even need a capital expenditure to put an OC48 on to the dark fiber. In other words I have the routers anyway and it costs me no extra to go from one building downtown to another building downtown on OC48s. Pricing for an OC48 the would cost me more than the price of getting dark fiber on my own to get a lambda from which I could get such a circuit just doesn't make sense.

***COOK Report***: It sounds like these companies, Yipes! and Telseon have been set up to do it in their metropolitan areas somewhat the same thing and that you want to do for your own customers?

**Freedman**: Absolutely. You are right. But it is not because I have grand dreams of becoming the "uber" telco. It is because some people want to buy something that is not available to them in the marketplace.

***COOK Report***: I also have the impression that these companies are set up with the intent of beginning to operate immediately on a very grand scale in a large number of cities and to go into the capital markets with stock offerings. Consequently all of a sudden there's a whole set of criteria that involves satisfying people other than your immediate customers.

**Freedman**: In my opinion the companies'combination of services and pricing did not make sense six months ago any more than they make sense today. To be fair I also must tell you that we have zero customers who are buying circuits from us today. That will be different in 30 to 60 days and frankly I hope there are many companies that will say "I will sell you an OC3 from New York to D.C. for $3000 a month".

***COOK Report***: You are welcoming competition along the lines of what you're doing. Is that what you're saying?

**Freedman**: Yes. Of course for a company that has zero circuits sold right now to welcome competition could be a pretty grandiose ego driven statement. Nevertheless my answer still is yes. And do you know what? I've talked to a couple of Metropolitan Ethernet plays. One of them actually said: "We have looked at our aggregated Ethernet product and based on what we learned I hope that all we use it for is out of band control to our boxes and that all we sell is circuits. Because of our software we can provision circuits quickly. We can do it economically." This person found out that when presented with a choice of purchasing either aggregated Ethernet or circuits, customers preferred circuits.

***COOK Report***: It sounds like you are saying that there is a significant perceptual disconnected between what Yipes! and Telseon think their markets to be and what people really want?

**Freedman**: Yes and quite frankly I think that companies like Enron will eat their lunch because Enron will soon be in dozens of buildings and in 50 cities. Enron will be in a position pretty quickly to say "You want to the OC3 from there to there. I will turn it up for you in 15 minutes and sell it to you for as long as you like. If you want a ten year contract, I will do the financial derivatives (or price protection) for you. And when you are done with that I'll sell it to someone else."

Net Access is obviously not going to dominate this market. I don't know for

sure but I do expect Enron to enter this business. Now therefore it could be that some point in the future Yipes! and Telseon will also figure that out and start selling circuits instead of aggregated gigabit Ethernet.

***COOK Report***: What do you hear from XO which owns 25% of level 3's dark fiber?

**Freedman**: XO has a wavelength product which means that you are basically on GigE and unfortunately costs about as much as dark fiber. The only thing to be said in its favor is that apparently it is not aggregated. Still, XO sells circuits and from that point of view as far as I'm concerned they are a telco. And again my question becomes "Am I going to trust the telco when they tell me that my wavelength is not aggregated?" The question is what does the market want? The people whom we've talked to would rather have an OC12 than a Gig E. That is strange you might say because an OC12 is less bandwidth and costs more to put IP on.

***COOK Report***: In other words they know they have the bandwidth.

**Freedman**: Well, they are 99 percent sure of it. There are strange devices that let you do OC12 emulated over IP which doesn't work and is a completely ridiculous idea. For all intents and purposes they can be sure however that they have an OC12. Whereas with the Gig E they have to be a little bit nervous and with the Gig E that someone tells them is aggregated they must be somewhat more nervous - somewhere between somewhat and a lot depending on what one knows about the company.

My conclusions here have nothing to do with the technical quality of the company because these companies have very smart people. They don't have anything to do with funding because they may run out of the money at the end of the year or they may not. But the same is true for many companies.

***COOK Report***: What about Cogent?

**Freedman**: The issue with Cogent is that it can't be done.

***COOK Report***: And despite the fact it can be done, they have gotten Cisco to invest several hundred million dollars in them?

**Freedman**: Why they would do that is certainly a mystery to me. My calculations show that the gross cost of providing what they want to sell is more than what it would cost them to run the network for the next three years. Remember an OC-192 will allow you to do 10 gigabit Ethernets — they are selling gigabit Ethernet to ISPs for VPN purposes for $10,000 a month per location. Now if they were selling $10 per megabit circuits no problem. But as soon as you have to take bandwidth and make it IP they have a problem. The reason for this is that if you want to make a redundant network with ingress and egress in order to step up and step down, you will have to pay an order of magnitude more just for the equipment to provide the service to the ISP. Now it is true they may not use an entire GigE if they are aggregating it, but he who sells to and ISP and plans on the full amount not being used is a fool.

If I were Cogent customer, which I may well be soon, I would light (fill) that GigE up because you know what happens to bandwidth if you don't use it. It goes away. If I don't use the GigE it gets aggregated. If I do use it they can't aggregate me without affecting other customers. And on the IP side if I wanted to take Akamai's traffic and ramp it onto Cogent, we'd bankrupt them because they pay transit to get to UUNET. The only question in my mind is how long are they going to be able sell me at $20 a megabit bandwidth that they pay $200 a megabit for. They are hoping they can get peering but the fact is that if people don't want to cooperate in putting circuits in they are either going have lousy service or have to send the bandwidth through transit which will cost them a huge amount of money — namely about 10 times as much as they're going to be collecting.

## Cogent and the OC-192 Brick Wall

***COOK Report***: So is part of the problem with Cogent an issue of scale?

**Freedman**: Yes. IP bandwidth does not scale right now. It costs more to move an OC-192 than it does to do the same for four OC48s. It used to be that it cost less to go up the food chain in IP. The broader your bandwidth, the less it would cost to send a single bit. That stopped happening somewhere between the OC12 and OC48 levels about a year ago. Let me explain to you the two reasons why OC-192 ports are expensive. One is that lasers capable of doing 10 gigabits per second are in limited supply. Furthermore, supplies may become more limited and not less so.

Look at the OC-192 card. Not only does it have an expensive laser on it but it also has tens of nanoseconds to move a packet. The refresh time for the memory in your computer is on the order of 30 to 60 nanoseconds. The refresh time on the very fastest RAM made is six to eight nanoseconds normally. This would be the cache RAM that is off the CPU but near it. I am not a hardware guy so I don't know what the access time of the ram on the chip is. But you certainly can't put a gigabit of RAM on a chip without it being unusually expensive.

Remember that if you have packets coming in at the rate of 10 gigabits per second and you want to buffer it for one-tenth of a second, you need, all of a sudden, a lot of RAM. If you want to buffer it for an entire second you need 10 gigabits of storage where you have to index and store packets somehow because of course they're not all the same length. There are a lot of bits on which you are going to have to do something like a two nanosecond lookup. I grant that you can do a large number operations in 10 nanoseconds on a modern computer. But only on stuff that is on the chip. Maybe I could see how to do an IPv4 look up in 10 nanoseconds. But if I need to do anything else like an IPv6 look up, it is a lot more bits that I simply do not have time to grab. If I have to keep everything lo-

cally right on the CPU, I have a very expensive computer that I must make in order to do OC-192 routing.

Remember that there are different philosophies of routers. On the backplane you have something like the Juniper. Avici believes that you should need a PhD to read the manual. Nortel believes they have come up with a simpler way except the only problem is they haven't yet made the chips that will implement the simpler way.

Consider then the backplane of the Juniper M160 which can do eight OC-192s. That router is roughly $300,000. The M5 is their smallest router and has the same switching capacity as an M40. So the M5 is perhaps one quarter of the switching capacity of an M160 and costs some $40,000. The M160 is basically four M40's glued together.

Now Cisco has been very happy about possibility that if you have a Juniper OC-192 network with one flow on it (that means nothing but communication between to post on the same source port) that it could possibly reorder packets. This is the worst that Cisco could say. Of course Juniper could say "Go read the Cisco Release Notes. They crash all the time and their software doesn't work." The hardware is expensive to make but it does a lot. The Juniper chips each have to move a quarter of eight OC-192s.

Let's say that I really want to do OC-192. I will have to have a $300,000 router and a $250,000 OC-192 card on each end. Frankly I really need another $300,000 router in front of it to do OC-192 to Gig E because there's no magic box that translates OC-192 to GigE. There are switches that will do this but they are also very expensive. So it is going to cost me a fair amount of money to turn IP into OC-192 to Gig E. If you add up the parts it's going to cost probably million dollars an end to be able to stick the aggregation stuff in. What I am doing is taking a 10 gigabit Ethernet and throwing it into a router with some redundancy over an OC-192 that's going to go somewhere. Furthermore, when I do this I still have not paid for the cost of the fiber equipment.

***COOK Report***: Sycamore is talking about switching of wavelengths or lambdas. Are they going to be able to take it the expensive scaling issues that you have just explained at the IP level and push them down one level in the protocol stack with any greater cost effectiveness?

## Switching MPLS Ciruits

**Freedman**: The switching things that people are doing seem to be: I can do faster that which you can already do with your dumb optical gear. But using dumb optical gear actually makes me happier because then there is no software complexity and software is what makes things crash that otherwise wouldn't. Or the switching people say that they are building a router in to their switch. But then they have the same problem that the router makers do. Their solutions seem to be little cheaper than just putting the optical gear and router next to each other. They still have the problem of moving packets in 10 nanoseconds.

There are things that people are playing with as, for example, at Nortel where they have queues that are made up of fiber loops and that they can use to inspect the packet before it gets there and in a microsecond to instruct a mirror to redirect the packet from the input to one of eight queues and have the queues all go to different interfaces. And that is what people are playing with. It is interesting in a sort of MPLS sense. All in all I am not aware of anything that is more cost-effective now than taking a Juniper M160 and mating it with fiber gear if you want to do IP. However, if you want to do MPLS, in other words if you want to do just circuits, then MPLS is much cheaper than doing routed IP (in the core, at least).

***COOK Report***: So to the extent that people are enthusiastic about MPLS it may be because it handles these scaling issues?

**Freedman**: That is for sure. That is what people are interested in. But you know what? You still have to get the IP on to the thing at the end. There is no 10 Gig E MPLS thing that takes layer two and moves it to layer three or takes layer three traffic and moves it onto MPLS. With MPLS we still have to have an IPto MPLS "thing" in front with an OC-192 card inside that it speaks MPLS and has to be able to handle OC-192 IPin order to get it on to the network. At least you have to have this until the point that you just say to someone "I'm going to give you only an MPLS circuit from here to there".

***COOK Report***: My understanding is that a part of what CANARIE wants to do with the OBGP protocol is to give Net Access like customers the opportunity to interconnect wavelengths that exchange points independently of carriers and to keep the bandwidth of such wavelengths off of the major transit Internet backbones. I gather that you are saying there is good reason to do this?

**Freedman**: Absolutely. There certainly is. Let's look at it as an analogy to MPLS. The question is do you want to build a just and MPLS infrastructure and hand people an MPLS pipe and tell them it to provision it somehow? I'd rather have someone say "Here's your circuit Mr." He may do this by going into the command line and provisioning my MPLS circuit from one point to another. Or a software package to use in the way that people who provision ATM circuits at Mae East now have. If someone wants to do an MPLS to MPLS interconnect, doing so is not as stupid as doing ATM. You least would not have the data loss that comes with ATM. So overall economically it might be workable thing. But I don't think that doing it by using software to automatically tear things up or break things down is a particularly good idea because the software that we have seen that would do this is not of good quality in design or implementation.

***COOK Report***: This sounds like just one more interesting example of Internet issues related to problems of scaling — problems that I'm hearing about more and more often.

**Freedman**: Let me say that such scaling

issues are not presently a problem for the growth of the Internet as we know it. The growth scaling bottleneck comes from networks that don't want to peer with each other. Internet growth has slowed down because broadband isn't and networks don't peer with each other. The fact that there are scaling issues in going to the next level of bandwidth doesn't matter yet. We are not some bumping into them yet because your average DSLAM has 500 users at one megabit and on the other end is a 34 megabit ATM pipe. Your average cable plant has 200 ten meg users on a 10 megabit uplink that does not have adequate bandwidth fed into it. This is what I mean when I say that broadband isn't yet.

And then you go to the other hand and say I'm going to hook up my fancy 10 gigabit Ethernet pipe to provider "x" and expect 10 gig to everywhere. You must ask to whom does provider "x" send those alleged at 10 gigabits. In fact if provider "x" has 40 peers to whom does it have it even as much as one gig Ethernet available? Not many. So there are utility problems in the infrastructure of the Internet that have to do with politics and poor design decisions at the last mile. In my opinion, these issues are making it so that we are just not seeing the problems that exist due to lack of scaling at the core.

***COOK Report***: Is there something inherent in the structure of circuit switched networks that has been keeping people from seeing the scaling issues that are now lurking at the edges of the connectionless Internet ?

**Freedman**: Rather than discuss issues of philosophy let me ask a rhetorical question. Why did the Y2K remediation happen successfully? Because it had to. There was no option. You would not have a job (or worse) if you didn't do it right. What was done properly was simply amazing. I will answer your question technically but I will also say that I believe that the reason the PSTN has worked is because it had to work. I guarantee you that the guy who maintains the 5E switch knows how it is used. I will also guarantee you that most router vendors (except for Cisco and Juniper) have not a clue about how they are used.

***Cook Report***: So the PSTN has grown up more gradually but also with a very firm sense of its purpose in comparison to the explosive growth of the Internet?

**Freedman**: I would say it somewhat differently. The PSTN had to work or you would lose your job. Number 2: It was simpler. Number 3: The vendors who made the gear understood more about how it was going to be used.

## Current Problems More of Stability than Scaling

So there are some technology issues here. It is more simple to do circuit switching and than it is to do packet switching in a reliable way. But there's also no excuse for the fact that there is not a router out there today that will happily move packets for ever and is hardware redundant. The main reason that Juniper is more stable than Cisco is because they only have a single forwarding path. They have a single point of failure in a sense in their smaller routers and are using a base OS that has big security problems. So although there are problems with the Juniper approach, people put up with it because the Junipers do not crash as often as Ciscos. Nortel actually has a clue but they don't have the hardware to put it on. No one else whom I have seen is actually doing a good job with software.

And there is no one who is actually doing something where you load your BGP module and run it on eight CPUs, you have the results come out, you pull out a board and the other guys keep running yet. You put in a new board. You put in a new BGP module and start getting updates and then once you see that it is doing the right thing turn off the other BGP modules.

I believe that there are four major problems with the Internet in terms of stability. Nortel actually articulated this for me and I believe that they're right. Router software. Router user interface that causes misconfiguration. Fiber. And a distant last is hardware. A router crash is considered a software problem. Outages that you see are not caused by routers physically dying. Hardware itself is not a big issue. Cisco routers and the 5E switch a been mature for about the same amount time. Yet the 5E is stable and Cisco is not because it still doesn't have a real operating system. In my opinion there really is is no excuse for the telco industry to be able to build switches today that are more stable than Internet routers. There are problems on the Internet but having routers crash because of stupid bugs that have not been rooted out of Cisco's operating system - which should not be a valid excuse.

***COOK Report***: Is Cisco's enormous growth a reason why?

**Freedman**: I think that the reason Cisco doesn't have a new operating system now is that people got fed up and left and went to Juniper. Cisco didn't have enough people to do a new operating system.

Let me state the problem another way. Networks have not demanded stability above all else. If networks said: "I will not buy your product if it crashes," we would have more stable routers. In fact the issue of stability is the reason why Juniper has grabbed a huge part of Cisco's market share. Five years ago Cisco had 90% of the provider market and that was five to ten percent of their revenue. My current estimate, while not based on any formal information, is that about half the packets that travel the Internet will hit a Juniper because an amazing number of networks are using them in their cores. Probably half of the "Cisco powered networks" are using Juniper and Cisco has been forced to turn a blind eye to these defections.

***COOK Report***: So aside from stability what you see as the key issues underlying Internet growth over the next couple of years?

**Freedman**: Until broadband really becomes broadband and networks get better and bigger connectivity to each other I think we are still at least a year way from

having scaling problems all over the place. UUNET is having scaling problems that they are really the only ones having today.

**UPDATE**: Since our interview Netaxs has sold 4 dedicated circuits on their Philly ring, and 7 circuits on the NYC-DC route, though the system is being brought up still, so the customers aren't live.

## Response from Bob Klessig

Mr. **Freedman** makes a number of negative comments about Metropolitan Area Network providers, such as Telseon, who use Gigabit Ethernet in their infrastructure. It appears that many of his concerns are based on the fact that he is considering how a service such as Telseon's would work for something for which it is not well suited. In particular, he is looking to connect two large routers together with a Gbps or faster link. It is safe to assume that the traffic between the routers is highly concentrated and that the routers will be in place for a long time. As Mr. **Freedman** says, this is a great scenario for a dedicated link such as dark fiber if it is available.

But Telseon is offering a switched service. It provides for new data connections in minutes to hours between any two points on a Telseon metropolitan network. Bandwidth choices are highly granular and quickly changed. The Telseon network is designed for businesses that need to have dynamic relationships with customers and partners.

Certainly the communications business tends to hyperbole with a strong tendency to characterize every new idea and technology as the Glorious Ultimate Solution for all known communication problems. At Telseon we avoid such claims and understand that a switched service is not the absolute best fit for Mr. **Freedman**'s need to connect big routers. But describing our service in the wrong context obscures information that is useful to the readers of the *Cook Report*.

What is the service that Telseon provides? It can be summarized by the following points:

It is a virtual connection (we call it a logical wire) with an Ethernet interface (10/100/1000).

Each logical wire has a Bandwidth Profile that allows bursts up to 75K bytes and a selectable average bandwidth or 1, 2, …, 20, 25, 30, …, 100, 150, … Mbps. The price of the logical wire is a function of the Bandwidth Profile with the cost per Mbps dropping as the bandwidth increases. By using Gigabit Ethernet switches, link aggregation, and wave division multiplexing, our network equipment costs are about one tenth that of a similar capacity SONET ring network.

New logical wires between two authorizing customers can be set up in minutes via our on line provisioning system and in a few hours via a telephone call to our Customer Support Center. Bandwidth Profiles on existing virtual connections can be changed even more quickly.

Whenever a logical wire is established or the bandwidth profile increased, Telseon developed software analyzes the network to insure that both a primary and a secondary path exist with sufficient spare bandwidth to fully accommodate the change. If capacity does not exist, the change is not effected. If capacity does exist, the bandwidth is reserved. We do not oversubscribe our network.

Utilization reports on a logical wire are available via a standard Web browser.

Primary and secondary paths are always available for each logical wire to fully protect against a single link or switch failure.

Bandwidth Profiles are enforced at the ingress to the network to ensure that one logical wire cannot interfere with another logical wire.

Security features include intrusion detection, MAC address spoofing lockout, and monitoring for unauthorized customer equipment on logical wires.

The logical wire operates at layer 2 that means that it is transparent to IP addressing schemes and the IP routing protocols. Our Service Level Agreement provides for 95$^{th}$ percentile one-way latency of 5 ms and packet loss of less than .004%. The SLA also provides monetary guarantees for downtime and missed installation commitments.

We could go on but we think that this makes the point that there is more to our service than just throwing together some Ethernet switches. Furthermore, we agree that ultimately it is our performance that will earn us the trust of our customers. We would welcome the opportunity to go into the additional details of our network with Mr. **Freedman** to start to earn his trust. Several of the principals in Telseon have experience with industry standards and the attendant "mumbo jumbo." We know it when we see it and you have our "no mumbo jumbo" promise!

## Avi Freedman Responds

Just a few points on Bob's response:

There are no applications or software demanding highly granular and changeable bandwidth today that I know of from Net Access or Akamai.

The base cost of Telseon and Yipes! 1mb-over-100mb/1000mb services are as much as customers will shortly be able to pay Net Access, and probably Enron and others for a 100mb provably dedicated pipe.

Analyses and promises are no substitute for proof of infrastructure where the possibility for aggregation follows.

SLAs are interesting if you'll guarantee me 3-4 months credit for any failure due to aggregation. An SLA can cover 99.9999999% but if the financial teeth aren't there, it's meaningless.

And in summary, I wish them well, of course - I just don't see the application for their model in the enterprise or ISP markets.

# ICANN and VeriSign in Alliance to Reinstate Defacto NSI Dot Com Monopoly in Return For Financial Support of ICANN

## ICANN's Pattern of Fraudulent and Deceitful Action Continues With Board's Capitulation to Staff and Vint Cerf in Melbourne

ICANN was imposed on the Internet with the promise that it would limit itself to technical coordination and abjure policy. It was created with the promise that it would bring the arrogant behavior of Network Solutions under control and free the Internet from having to operate under the growing dominance of a company uninterested in playing by the basic rules of internetworking norms. And most of all it was hyped as the fulfillment of the industrys promise to self regulate so that the US government would not have to. It is now starkly clear that all these promises should never have been taken seriously.

ICANN's JDRPlawyers run it according to their own agenda which, to the extent it can be fathomed, is that of the trademark industry. These attorneys, and the staff speak a language that purports to make decisions based on cooperation and consensus spawned by open decision making. The words of this language bear no relationship to any reality. Listening to ICANN is like listening to the cascade of lies purporting to be truth that comes flooding forth from Big Brother in Orwell's 1984. The arrogant clique of ICANNites gets away with what they do because they have established a rogue regulatory bureaucracy so complex that only a handful of full time ICANN watchers can adequately grasp the enormity of the fraud being perpetrated in the name of the new Board Chair, Internet "father" Vinton Cerf. Hash words are these. But we believe they are deserved because the duplicity of these people seems to know no boundaries.

What follows is the latest example. ICANN is in full collusion with Verisign as the new owners of Network Solutions to run the Internet through the DNS control point as a duopoly. The public has been fed a myth that ICANN was established to carry out a mission of technical coordination and consensus management designed to allow the internet industry to self regulate. This statement of ICANN's "purpose" has been fraudulent from day one. What we have been given instead is an ICANN -Verisign duopoly that is also free from any of the constraints of due process to be expected in the process of government regulation based on civil and administrative law instead of on the whim of two parties. The one motivated by commercial greed while the other is motivated by reliance on the first as the predominant source of its revenue.

Lets examine what happened. From http://www.icann.com/melbourne/proposed-verisign-agreements-topic.htm we read: "Perhaps most relevantly, VeriSign's once-dominant market position has been severely eroded. VeriSign's share of total registrations has fallen to about 50%, its share of new registrations to under 40%, and its share of net new registrations (taking into account non-renewals and transfers) to an even lower level. This trend appears to be continuing in 2001.

"For all these reasons, when ICANN and Verisign began to discuss VeriSign's plans to divest itself of its registrar business so as to qualify for the automatic four-year extension to operate the .com/.net/.org registries, it quickly became apparent that the importance and value of the separation of ownership of VeriSign's registry and registrar businesses to ICANN and the community had diminished quite significantly over the 15 months since the original registry agreement was signed. While VeriSign might well wish to retain its registrar business, the fact that separation of ownership will automatically extend its ability to operate the .com/.net/.org registries for an additional four years is a powerful incentive to cause that separation to happen. On the other hand, that ownership separation is clearly not as valuable to the community or ICANN under today's market conditions as it appeared it would be at the time the agreement was signed."

"Given these circumstances, the management of ICANN and VeriSign began exploring whether there was an alternative set of arrangements that would be more attractive to both parties. The result of those discussions, which have been ongoing since last summer but more intensely over the last two months, is a proposal that VeriSign has now made to the ICANN Board to amend the existing registry agreement."

What the obscenity of industry self-regulation "ICANN style" has permitted is the unchecked growth of the alliance between ICANN and Network Solutions the first stage of which was brokered by IBM in the secret meeting held in down town Washington on July 31, 1999. The ICANN strategists at IBM and ISOC had ridden to power trumpeting a holy mission to restrain an unfettered greedy monopoly given NSI by the National Science Foundation. Becky Burr, as a matter of alleged public interest, had wrested control of the government's power over NSI from NSF and given it to the IBM friendly Commerce Department. Magaziner was able to give birth to a highly secretive and authoritarian ICANN because he, Mike Roberts, Esther Dyson, John Patrick and others who were involved were shrewd enough to focus attention on the perception that NSI was the evil monopoly from the grasp of which ICANN would free Internet users by opening the domain name space and

introducing competition. I asked Magaziner in a September 1998 interview what got the US government involved in the DNS, He told me that in November or December of 1996 representatives of two large corporations asserted to a meeting of his Internet task for that the commercial promise of the Internet would never be realized if trademark violations in DNS names were not brought under control. This is what powered well hidden the moves over the next two years to form ICANN.

The problem by the summer of 1999 was that no one had stepped forward to pay ICANN's bills. ICANN's request for an annual one dollar per domain name tax had been met with derision. ICANN's assertion of control over the DNS and trade mark issues depended on Network Solutions becoming an accredited registrar. Burr had been threatening to recompete the Network Solutions contract. NSI's long term financial interests depended on postponing such an event for as long as possible. ICANN's long term interests depended on tapping into the NSI cash cow in order to be able to collect enough money to pay for the multi million dollar annual bills of its law firms, business class tickets for board members quarterly junkets to board meetings and so called "essential tasks" like the $440,000 bill announced on March 9, 2001 http://www.atlargestudy.org/budget7.shtml for Swedish technocrat Carl Bildt's study expected to justify getting rid of the at large membership.

The result was that at the end of September of 1999 ICANN and NSI signed agreements that turned them from enemies into the best of friends. In return for guaranteed payment of a major portion of ICANN's budget NSI was guaranteed control of dot com for several additional years. Given the lack of any oversight NSI and ICANN remained fee to continue their collusion. ICANN remained in need of guaranteed budgetary cash flow. NSI's new masters (Verisign) wanted their monopoly on dot com reinstated and guaranteed in perpetuity. Never mind that such action was a stark reversal of the commitments on which ICANN grabbed what power it now holds.

Here is how Verisign CEO Stratton Sclavos justified his actions in a letter of February 28, 2001. http://www.icann.com/nsi/sclavos-letter-28feb01.htm

"For all of these changes in the term periods and provisions in the agreements for .com, .net and .org, and the further commitments discussed below, we have mutually agreed to modify the current provision in the .com agreement concerning the procedures for any subsequent agreement."

"We have accepted new and substantial obligations on the part of all three registries to pay fees, as part of ICANN's cost recovery program, identical to those paid by similarly situated registries. Because there is an existing fee structure already in place for these registries and the accredited registrars that register names in them, there will have to be a transition plan to move to the new fee system; that transition will be the subject of continuing negotiations between ICANN and VeriSign, with the objective of coming to an agreed transition before these agreements become final that can be put in place by the end of ICANN's fiscal year on June 30, 2001. ICANN has committed that the final agreement will ensure that VeriSign will not pay more than its proportional share of the overall Variable Registry-Level Fee (currently capped at $3.5 million for all registries) as measured by its share of registrations, and that an appropriate mechanism will be included to ensure that any future increases in this Fee can be offset by appropriate increases in the fees that VeriSign charges to accredited registrars, on the same basis as that set forth in the registry agreements for new TLDs. We believe that this regularization of the fee structure with respect to the .com, .net, and .org registries will further assure ICANN's financial stability and enable it to accomplish the important tasks ahead."

ICANN and Verisign had in effect conspired to ensure ICANN's financial viability by building in the right to arbitrary annual rises in fees. In a bizarre twist Verisign had to be made to appear to be giving up something. That something turned out to be the dot org registry. In a move clothed as a gift to non profits and NGOs, they falsified the purpose of the dot org name as being always intended for non profits and promised to find a 'non profit' organization to run the registry and return it to its original purposes — adding that such unnamed organization would determine the conditions under which registrants could keep their names. In doing so they delivered to the trade mark industry the opportunity to put out of business dot org web sites set up to criticize corporations – such sites often being named *"corporation"*sucks.org

As Sclavos phrased it: "We would agree to terminate our operation of the registry for .org on December 31, 2002 and to cooperate with ICANN in transitioning .org to management by a new, non-profit organization representing the global universe of non-profit organizations. Among the issues to be determined in this transition is whether .org should be limited to registrations only by non-commercial entities, and if so, what transition arrangements need to be established for those existing registrants that would not qualify under that limitation. ICANN has agreed that, at a minimum, existing registrants would be permitted to remain in the new .org registry for one renewal cycle under its new management. In addition, and as another part of the transition process, all ICANN-accredited registrars would continue to be permitted to register qualifying names in the .org TLD for three years after termination of our operation of the .org registry, during which period the new registry could develop whatever registration policies for the future it thought appropriate. Our objective is to provide a permanent and affordable home on the Internet for the non-profit sector and in so doing make a major effort to close the digital divide on a worldwide scale. To see this through, we also would agree to support the new non-profit .org registry operator, and provide it with a contribution of $5 million that would be used for an endowment to help cover its operating expenses. Further, we also agree to make available to the party designated by ICANN as successor operator of the .org registry the use of global resolution and distribution facilities at no charge for one

year."

Of all the violence done by language to the truth in these agreements Sclavos'assertion that "Our objective is to provide a permanent and affordable home on the Internet for the non-profit sector and in so doing make a major effort to close the digital divide on a worldwide scale." — stands preeminent. What is diving these events is raw and naked corporate greed – nothing more and nothing less.

The Melbourne Australia ICANN Board meeting (March 10 -12, 2001) offered a predictable charade. Even the loyal opposition of long time watchers and critics found it difficult to contain themselves in the face of policy made in secret by Vint Cerf as new Chair, Jones Day Reavis and Pogue.

The not for attribution comments by key figures were filled with outrage. Critc: "With no warning, with the text of the agreements not even published, the Board accepted a series of contracts that turn ICANN into a super-regulator for the internet. These new rules will make ICANN rich, and make ICANN a highly intrusive regulator of registries. And this happens without advance notice to the public that this decision was to be proposed. Indeed Karl didn't know it was coming. I bet they didn't tell Andrew M-M either. I think we've all just been had, big time." Observer: "You guys still think you can reform ICANN?" Critc: "No. Right now I think the time has come to ask Congress to abolish it. Maybe I'll feel calmer after a good night's sleep."

# BWG and Other Mail List Reaction

On March 1, **Ellen Rony** commented to BWG.

ICANN writes, " the management of ICANN and VeriSign began exploring whether there was an alternative set of arrangements that would be more attractive to both parties".

**Rony**: Why is a technical coordinator working on arrangements that would be [financially] attractive specifically to VRSN?

"ICANN management believes that there is a persuasive argument that amending the existing registry agreement with VeriSign as proposed would be of far more benefit to the Internet community, and do more to enhance long-term competition, than would the continuation of the existing agreement"

**Rony**: ICANN management? That be whom, Sims/Touton/McL?? Persuasive argument? Let's hear it.

"The rationale is that ownership separation is no longer necessary or useful in promoting competition, so long as the structural separation is effective in accomplishing the basic purpose."

**Rony**: Could someone please translate this ICANNspeak for me. Separate buildings required, but not separate ownership? Buildings provide a separation of conflict but financial assimulation doesn't?

"This reflects ICANN's belief that there is little if any additional competitive value under today's market circumstances in forbidding the registry operator from also being a registrar, so long as it is done is such a way so as not to discriminate against other competitive registrars."

**Rony**: ICANN's justification for just doing a 180 degree turnabout.

"VeriSign would agree to permit any ICANN-accredited registry operator (including .org) access to its global zone resolution and distribution facilities at terms to be determined. "

**Rony**: Would someone please explain the technical ramifications of this.

"The elimination of special rules or provisions dealing with VeriSign is an important step forward in the ICANN process."

**Rony**: All the while making special rules or provisions in dealing with VRSN.

"ICANN has agreed that, at a minimum, existing registrants would be permitted to remain in the new .org registry for one renewal cycle under its new management."

**Rony**: Meanwhile, registrars have been collecting fees for playing upon the paranoia of registrants to protect their names in all three TLDs. Does ICANN really need to give us yet another reason to hate it's manipulative ways.

Also on March 1 **Milton Mueller** wrote: I just read the contract. This is regulation by contract. Prices are set, reservations are made to set policy (as Michael points out below), not just technical standards but codes of conduct and particular data formats. It's also a way to finance ICANN. UP to15% annual increases in the contribution of the registry industry to ICANN, for no particular reason.

Isn't this the template contract they will apply to the new TLDs? I'm afraid my longstanding FCC analogy is being played out to the max, to a degree that surprises even me. I see that Froomkin as usual is already fully on top of this. He's right. This contract represents the fully developed model of what ICANN has/will become. Complete leveraging of the root to regulate the industry and impose name-related policies.

**Rony**: RFC1591 states: ORG - This domain is intended as the miscellaneous TLD for organizations that didn't fit anywhere else

So ICANN is unilaterally scrapping the RFC definition and providing a jackpot for VRSN, all the while providing insufficient time for public comment before completing this landmark agreement. Further, by indulging in such market manipulation as outlined in this agreement, it has blatantly violated its charter as a technical coordinator. VRSN's dropping market share in registrations was an inevitable consequence of introducing competition among registrars. While it may be at 50%, it still is five times larger than the second place registrar and even at the lower market share certainly has far more customers than it did two years ago. And what does any of this

have to do with technical administration of the Internet?

**Mueller**: the key here is to look at the ORG contract as a model for how ICANN plans to license all registries in the future: on a short tether, paying huge fees to ICANN, totally subject to its IP-friendly policies, regulated, etc. etc.

In the original concept of DNS, a name delegation, at the top or any other level, was a delegation of authority to manage a domain. This is a major shift in that structure. The administrator of the top level controls everything, and simply allows tld operators to operate name servers and collect some money for their troubles.

**Richard J. Sexton**: Roberts, Touton, McLaughlin say the new .org will be "for the specific use of nonprofit organizations," and that under this deal .org will be "returned . . . to its original function as a registry operated by and for nonprofit organizations." Go tell them that their history is bad.

**Karl Auerbach**: I notice that in these contracts the price that the registry may charge may be increased by amendment. But there is nothing to force the registry to reduce prices due to more efficient operations. In other words the contracts guarantee a price floor and do not form any sort of cost+profit allowance structure. So under these contracts there is no reason for internet users ever to expect a reduction in the $6 fee, NSI has a permanent $6 lock on .com names, forever.

**Mueller**: This is what happens when the people who regulate an industry have no experience in regulation and run around denying that they are engaged in regulation.

**Judith Oppenheimer**: There are only two corporate entities to this deal: Verisign and ICANN. They are the corporate entities the deal is structured to benefit. And its the board of one of them whose vote will close the deal.

One only need observe how enthusiastically yesterday's "Verisign" press conference was championed by ICANN, with opening remarks made by Mike Roberts and the press precluded from participating. This was in fact a congratulatory announcement staged for the financial community, and it worked. But ICANN is hardly a neutral industry facilitator of a proposal between Verisign and the public.

Conspiracy? No. Brilliant business dealing on Verisign Inc.'s part? Yes. On ICANN Inc.'s part, too. And that's the problem.

## Semantic Double Speak by ICANN's Attorneys

On March 2, **Bret Fausett**: I recommend for your weekend reading pleasure a very interesting exchange between Milton and the forces of evil (Touton and Sims) on the NC List.

It's a thread that started with my post on Thursday to the Names Council ("Reasonable Opportunity for Comment"). At various times, we had responses from Roberts, Touton, and finally Sims. Milton steered the conversation in the right direction, focusing on the various "policy" issues imbedded within the proposed Verisign contracts. Sims takes the rather incredible position that anything in a contract is not a policy (yes, the distinction is apparently a matter of form, not substance).

All the posts are linked from here <http://www.dnso.org/clubpublic/council/Arc04/mail15.html>

And here <http://www.dnso.org/clubpublic/council/Arc04/mail16.html>

I know this was the thinking underlying the rejection of Karl's request for reconsideration last year, but the philosophy of staff has never been laid bare so completely. Kudos to Milton for pushing this.

**Louis Touton**: How do you come to the conclusion that the NC "must approve" the agreements? The DNSO/NC's role within ICANN is to make recommendations for new substantive policies. It does not have any role in approving new or revised agreements.

**Mueller**: Utter nonsense. As you know perfectly well, ICANN defines and implements policy through its contracts with registries and registrars. If the DNSO cannot have any role in approving these contracts, then it has no influence over policy.

Your decision that ORG must be assigned to a non-profit organization is a policy decision. Deciding that registrations in ORG will be limited to "non-profit organizations" is also a policy decision (one that, contrary to your assertions, has no basis in RFC 1591 or any other prior decision, and could have a substantial impact on current registrants in that space).

Your decision that the current level of market competition justifies allowing integrated ownership of COM registrar and registry is a huge policy decision. Personally I agree with it but on procedural rather than substantive grounds, it flies in the face of ICANN's process to say that registrar and registry and other impacted constituencies have no vote on it.

Fixing the price of a registry is a regulatory policy decision. Taxing the registry to support ICANN (at a rate that increases 15% a year, interestingly) is a policy decision. Don't you think registrar and registry constituencies, including the ccTLDs, have a right to review and approve those decisions?

The fact that these contracts are used as a template that will most likely be applied to all future registries also has long term policy implications. A policy of "presumptive renewal" for COM is, well, a policy. Frankly, Louis, I am appalled at the ICANN staff's total abandonment of the bottom-up concept.

And on March 4: **Touton**: Thanks for transmitting the Gomes memo. It adds value to the discussion. To me, it proves that a major policy change (IANA being the [informal] policy authority at the time) took place in early 1996.

I could quibble with your interpretation of 1591 and 920 - as you could with mine - but isn't that irrelevant at this point? Whatever was the "original intent" in 1984 when ORG was created, there were

less than 100 registrations in it then. 5 years have passed and 2 million new registrations have taken place since the policy was changed, completely transforming the nature of ORG. We are facing a completely different situation now.

Changing ORG to turn it over to a non-profit, and restricting registrations in that space, may or may not be a good idea. But semantic debates about "original intent" don't contribute anything of value to that policy debate. Can we agree on that? It seems so obvious to me.

I have another important question for you. I am rather puzzled by your participation in this debate in a way that reveals a strong commitment to the policy that ORG should be run in a certain way. I am new to the Names Council, so forgive me if this is an uninformed question, but I was under the impression that domain name policies are supposed to be made by Board members, constituency members and their elected representatives. My understanding is that you are hired by ICANN as a staff lawyer to carry out the corporation's policies, not to make them. I certainly value and welcome your legal advice, but I don't think it is appropriate for you to be actively promoting a specific policy approach. Am I incorrect?

Date: Mon, 27 Mar 2000
From: "Gomes, Chuck"
Subject: Re: network solutions new pitch. X-
To: Michael Sondow

The topic of RFC 1591 guidelines for .NET and .ORG registrations was talked about many months ago, but it might be helpful to restate some points in that regard. In the first part of 1996, NSI was still attempting to enforce the RFC 1591 guidelines with regard to .NET and .ORG second-level domain name registrations. Unfortunately, checking for compliance had to be done manually, thereby making the response time much slower. This was compounded by the rapid growth of .NETand .ORG registrations. In addition, we found this: in attempting to ensure compliance with the guidelines, we found ourselves penalizing those who were honest and rewarding those who were willing to lie. If someone said they were an Internet service provider, we took their word for it, understanding that the definition of an Internet service provider was becoming increasingly blurred and that it would be extremely difficult to investigate claims made by applicants. A similar problem existed with .ORG, determining whether or not an organization was not-for-profit.

In light of these issues, we consulted with the IANA (Jon Postel). Jon specifically recommended that we stop screening for compliance and instead rely on registrants to choose the appropriate TLD. We then followed Jon's recommendation.

The point here is not to blame the change on Jon but rather to communicate that it was not a unilateral decision by NSI. The decision was made in consultation with the IANAand was made to deal with specific problems experienced when implementing the RFC 1591 guidelines. As others have pointed out, the Internet has changed drastically since RFC 1591 was written. The changes related to .NET and .ORG domain name registrations are eamples of dealing with some of those changes.

To date, NSI still attempts to enforce RFC 1591 guidelines with regard to .EDU registrations. That has become increasingly difficult as the number of registrations have grown. Moreover, registrants are typically very dissatisfied with the response times because the requests have to be processed manually. Fortunately, the number of .EDU second-level registrations still does not come close to the number of .NET and .ORG registrations in 1996.

**Bret Fausett**: Good ol' Andrew. Now he's an antitrust lawyer, too.

See: <http://www.centr.org/meetings/ga-9/legal-report.html>

**Jonathan Weinberg**: This is pretty amusing. The document in question is CENTR's draft ICANN-ccTLD contract. CENTR wants to call it a "contract for services" to emphasize its view that ICANN's role vis-a-vis the ccTLDs is simply to provide them with root-server services. Andrew answers that for CENTR to draft up, and ICANN to agree on, a single "contract" to be used by all of the ccTLDs to purchase services from ICANN smacks of illegal price-fixing or something, *so he tells them to change the name*. Silliness. Either there's an antitrust issue here or there isn't (I think there isn't), but if there is one changing the name isn't going to fix it.

**Bret Fausett**: Isn't one significant problem facing ICANN the fact that the original NSI-ICANN-DoC agreements did not specify the nature of the divestiture. If you listened to the Verisign conference call with analysts last month, you heard a number of "creative" solutions discussed about what the divestiture of the registrar would look like.

One thing they considered was "selling" the registrar business, but then contracting with the purchaser to continue to provide the "Network Solutions" customer experience under contract, using the name trademarks, web site, etc. on a fee per domain name basis.

I don't pretend to know whether there are legal standards on what constitutes divestiture, but I'm sure Verisign has spent significant sums on legal analysis of that issue over the last year. When you look back at the original agreements, they contain no guidelines on what the divestiture should look like, who the purchaser could be, what approvals ICANN and the DoC had over the plan, etc. In other words, ICANN made a terrible deal.

Against some of the things that NSI wanted to do with the divestiture, ICANN had zero contractual leverage to push it in any different way. This new deal is a way out.

## The Foreign Language Domain Cash Cow

COOK Report: In the meantime one of the reasons for judging as false ICANN's expressed concern about NSI's declining dot com registration percentage is NSI's unchecked foray into foreign language domain names. The first few months of

its registration of non ascii domains at $200 a name yielded several hundred thousand new customers. A good example of the mischief being done is reflected in the following exchange.

On March 4 **Ellen Rony** cited a quote from the e-zine The Register "Meanwhile, we're gonna have to get round to registering www.leregistre.com, www.elregistro.com and dieregister.com." Rony asked: That can't be correct? Those domain names don't require multilingual translation because they all use the same alphabet set as English.

On March 5 **Ted Byfield** responded: yeah—whoever wrote that was joking, Ellen, but making a sharp point nonetheless. with these 'multiligualizations,' which are driven more by marketing than by engineering or an eye on coherent policy development, we can soon expect to see 'collisions' of that kind.

**Rony**: French, Spanish and German (and other languages among the 50+ new ones in the testbed) use predominantly the same alphabet with exception of some letters with accents grave and agui (?), umlauts and such.

**Byfield**: yes.

**Rony**: So, if I register croix.com, whether in English and croix.com in French, it would require the same ASCII characters. Under this testbed, will French language be translated under RACE, so the French word "croix" becomes bq—tdgxcbvjnt or something weird like that. Wouldn't people prefer to see the familiar croix.com to some bq—weirdness. Thus, a word in French, Spanish, German that doesn't use the umlaut or unusual characters wouldn't go through any RACE transformation. Is that correct?

**Byfield**: maybe this'll help: In addition to being a "french" word, "croix" is also an alphabetical string that can be expressed in ASCII—so you can think of it as already-multingual. it doesn't need any RACE-type nonsense.

"à-l-américaine" is a French phrase that *can't* be expres- sed in ASCII without losing the diacriticals—Americanizing it, as it were. (Can RACE handle apostrophes? I doubt it.) So, in order to represent it "in French" using RACE, it'd need to be a "bq—" string.

The inclusion of several languages with strong common roots in Latin will initiate a whole new class of UDRPproceedings: is "liberation.org" confusingly similar to "libération.org"? There are traditional orthographic hacks for getting around problems like this, for example, the representation of the german "ö" as "oe". So the government formerly known as the Haider regime in Austria, which has studied at the feet of Milosevic when it comes to cultural manipulations, might say that it wants fpo.org—or fpoe.org (ASCII) or fpö.org (ren- dered RACE) based on how it decides to represent the "O" in "Österreich."

This mess takes on new levels of complexity with every lan- guage included in the multilingual testbed. and ICANN, our alleged 'technnical coordinating body,'is holding valedic- tory teleconferences with the organization ('chinese walls' notwithstanding) is bringing DNS to nadirs no one dreamed of only a few months ago—well, except for the knucklehead Paul Hoffman who intimated what a f… up his pet RACE was when he wrote:

Note that a zone administrator might still choose to use "bq—" at the beginning of a host name part even if that part does not contain internationalized characters. Zone administrators SHOULD NOT create host part names that begin with "bq—" unless those names are post-converted names. Creating host part names that begin with "bq—" but that are not post-converted names may cause two distinct problems. Some display systems, after converting the post-converted name part back to an internationalized name part, might display the name parts in a possibly-confusing fashion to users. More seriously, some resolvers, after converting the post-converted name part back to an internationalized name part, might reject the host name if it contains illegal characters.[1]

...thereby providing a 'technical'basis for the regulation of

[1] http://search.ietf.org/internet-drafts/draft-ietf-idn-race-03.txt

**Editor**: Finally the pre Melbourne ICANN board meeting string off arrogant decrees ended with put downs to the loosers in what began to be referred to as the ICANN lottery for new gTLDs.

**Fausett**: Here they come folks <http://www.icann.org/committees/reconsideration/>

"Because we conclude the process was fair, and resulted in a rational conclusion that met the objectives of the exercise as announced at the beginning, there is no basis for reconsidering the Board's selections. Even if, for the purposes of argument, there were factual errors made, or there was confusion about various elements of a proposal, or each member of the Board did not fully understand all the details of some of the proposals, this would still not provide a compelling basis for reconsideration of the Board's conclusion. Given the uniqueness of this process, the inherent subjectivity of certain of the criteria involved, the inevitable difficulty of reaching consensus on a fact-intensive evaluation of many times the number of proposals that could possibly have been selected, and the limited objective of finding a small number of acceptable proposals for this initial proof of concept, it would not serve the interests of the community to essentially allow these decisions to be reargued on grounds over which, at best, reasonable people could differ. Given ICANN's stated goal in this process, which we believe has been met, there is no compelling reason to reconsider the Board's selections."

**Judith Oppenheimer**: "Since it was clear from the beginning that only a limited number would be selected, no applicant could have had any reason to believe that its application would definitely have been selected, and thus no one can reasonably claim that they were misled in any way by the process — which was in every respect fully open and transpar-

ent."

Its the icing on the cake - "open and transparent". ICANN-speak for Who-Asked-You,Go-F*ck-Yourself.

**BWG attorney**: Bottom line: your $50,000 bought you a lottery ticket and nothing more. Everyone got the same unfair treatment, so it doesn't matter if we got the facts wrong as we were involved in an exercise in which getting facts right was not a necessary element of making a fair decision.

Hmm. ICANN's position is that applicants had no right to any particular consideration other than treatment equally bad or good as other applicants.

Two more reactions: 1. If ICANN got facts about A wrong but not about B, did A really get the same treatment as B if both were subjected to the same standard of indifferent care? I don't think so. Certainly if this were an administrative proceeding that wouldn't hold water. (Again we see how much worse off parties are when limited to 'private' remedies; as a general rule, private parties are allowed to be much more arbitrary, and much more sloppy, tan agencies when making decisions. Here's one decision that doesn't look at all 'good enough for government work').

2. If ICANN's analysis is right, it was running a lottery. Does that require any special licenses under California law?

BWG observer: As if we ever doubted who is the source of policy at ICANN, these messages from Declan McCullagh's politech forum show us how "consensus" for the latest contractual turnabout was obtained.

**Mclaughlin**: Declan, We don't have any intention of kicking out existing domain name holders. The idea is to turn over management of .org to some appopriate organization/association/entity/whatever, which would then make decisions about .org registration policy. I'm always amazed by the amount of misreporting & hyperventilation about domain name stuff — this one's no exception.

**Joe Sims**: Declan, since I negotiated the proposed agreements on behalf of ICANN, I thought I should point out that this particular issue is a red herring generated by bad news reporting. The proposed agreements make it clear that what will happen to .org after VeriSign gives it up, and what (if anything) will happen to existing registrations, will be the subject of the normal consensus development process over the next year or so. To repeat: contrary to some news stories (and an early, incorrect SlashDot posting), there have not been (and could not be) any decisions made on either of these issues, since those decisions are obviously policy issues that must be resolved by community consensus. Obviously, even if future registrations in .org were limited to non-commercial organizations (which clearly was the original intent), existing registrations could easily be grandfathered if that was thought to be the fair and equitable thing to do, as many people argue. The point is that this should and will be debated to a community consensus, which is the way the process should work, so no one is in jeopardy now.

**Vint Cerf**: The proposal is just a proposal - and the question of the presence of non-not-for-profits in the .org registry is still quite open. ICANN has no desire to create hardship where there isn't any. It seems unlikely that existing registrants would be "evicted" without a good deal of discussion and planning and for the sake of simplicity and fairness, it would seem more reasonable to limit FUTURE registrations - the messy part is the possibility of some kind of gold rush to register before such registrations (ie of for-profit organizations) would no longer be accepted.

**Byfield**: That's hilarious: Mclaughlin, Sims, and Cerf all respond in unison? you know something's up—and it's not that the reporting was wrong. Witness the fact that Sims, who makes more hay about bad reporting than the others, yet again reasserts that .org had a positive 'ori- ginal intent.'

**Andrew McLaughlin**: We don't have any intention of kicking out existing domain name holders. The idea is to turn over management of .org to some appopriate organization/association/entity/whatever, which would then make decisions about .org registration policy.

**Byfield**: translation: we'll delegate the power to kick out current holders.

**McLaughlin:** I'm always amazed by the amount of misreporting & hyperventilation about domain name stuff — this one's no exception.

**Byfield**: (except when it's IP psychotics who are hyperventilating, natch.)

**Joe Sims**: Obviously, even if future registrations in .org were limited to non-commercial organizations (which clearly was the original intent), existing registrations could easily be grand fathered if that was thought to be the fair and equitable thing to do, as many people argue. The point is that this should and will be debated to a community consensus, which is the way the process should work, so no one is in jeopardy now.

**Byfield**: Jeopardy — say no more.

**Vint Cerf**: The proposal is just a proposal - and the question of the presence of non-not-for-profits in the .org registry is still quite open. <...> It seems unlikely that existing registrants would be "evicted" without a good deal of discussion and planning and for the sake of simplicity and fairness, it would seem more reasonable to limit FUTURE registrations - the messy part is the possibility of some kind of gold rush to register before such registrations (ie of for-profit organizations) would no longer be accepted.

**Byfield**: translation: there will be a good deal of discussion, then they'll be evicted—because, of course, 'registration'includes renewal. It's clear as day where this is headed.

# ICANN Watch Site Revamped Major Resource Created

In the second half of 1999 attorney David

Post, got together with Michael Froomkin to create a web site called ICANNWatch www.icannwatch.org . It was posited on the suggestion that ICANN would create the most significant impact imaginable on the future of the internet and therefore that what ICANN did and how it evolved was a matter of major concern and deserved "watching". Dave Farber joined as a third sponsor editor soon after.

The site served as a place for posting and then warehousing essays that Post and Froomkin wrote as well as essays by others. The essays were usually quite good. Unfortunately whatever contribution Farber made to the site was never clear. He himself contributed no essays.

Early this year Michael Froomkin began devote himself to a major reorganization of the site. The result has been open since the end of February. It is a remarkable achievement. More important an achievement that is quite unique. For the ICANN knowledgeable it serves as a kind of guide to what its happening and a gateway with live links to major resources. It is a long needed tool that must be applied to following what ICANN is doing for one of the strengths of those pushing ICANN has been that with out any balance, openness and accountability ICANN has been free to tell repeated lies and do so with impunity. What's more, ICANN has gotten away with doing this because only a tiny handful of people have had the motivation to devote the countless hours to acquire the necessary expertise.

ICANN's method of operation has been to talk out of one side of its mouth about how it does nothing besides very minor consensus based technical coordination of the DNS. The goal would seem to be one of convincing the press that it doesn't have to do more than publish ICANN's press releases.

The new site is designed to put on display as full a record of ICANN's operations as can be gleaned by outsiders. ICANN novices who visit and explore should be immediately struck by the depth and breadth of the activities chronicled. No innocent operation this. The activities of the hydra headed monster are now on full display.

The site is now centered around an edited web forum. At the top left of the home page the reader will find the topic "Inside ICANNWatch" with the following Links

LOGIN HERE
Home
Current Topics
Send News
Propose New Thread
Editors'login

Our Mission
ICANN for Beginners
FAQ
From the Archive
About us
Recommend us
Write to us

Under about us we read: "ICANNWatch.org is brought to you by the joint efforts of three different groups, each essential in its own way.

First, there are the selfless techies from Mindshare.net without whom none of this would be possible. They have donated the servers, plus installed and continue to maintain the code. (This site uses PHP-Nuke, which is free software released under the open source GNU/GPL license.) Second, there are the editors of ICANNWatch. The founding trio was David Post, Michael Froomkin, and David Farber, but we expect the list to change over time. You can find a current list and short bios here. [**Editor**: clicking on the link "here" brings one to the editors page which is much improved over the original site. There Dave Farber is described as "a founder editor of ICANNWatch," who "has not been active lately." Jonathan Weinberg is described as the current third active editor.] Third, and perhaps most important, are the readers and users of ICANNWatch, many of whom sent in content that we published, and feedback that we appreciated, back when we were a static site. One of the exciting things about our transition to our new dynamic format is that it will increase the ways in which our users can interact with us and (more importantly) with each other.

We hope that the discussions on this site will become a part of reaching consensus about where we go from here. We expect that these discussions will also aid in establishing when, and whether, there is a consensus for ICANN-related activities. We hope also that this forum will provide an occasion for us all to think broadly at times about how to ensure that this communicative tool we call the Internet remains free, accessible, and best able to contribute towards human welfare."

The three editors post short essays that they hope will become discussion threads contributed to by site users who are invited to register for a log in and password that will allow them to contribute to the threads. We note that we found the log in process mildly confusing. A conversation with Michael Froomkin revealed that to be an artifact of the PHP-Nuke software that the site is based on. This software apparently makes site administration much easier. Given the fact that the site is quite complex and supported entirely by volunteers, it was the only reasonable way to achieve the desired capabilities. At any rate as long as users accept cookies, they will only have to log in once with their newly assigned id and password. Given the rather slow response time of the site this is something to be generally grateful for.

Beneath the Inside ICANN Watch box is another box called Useful ICANN sites

**ICANN itself** [ICANN home page]
**Berkman archive** [Berkman Center site with archival data from earlier ICANN meetings]
**ICANN Blog** [Brett Fausett's web log of valuable pointers to other ICANN sources]
**Roving Reporter** [Ted Byfields icann commentary]
**ICANN News(DNH)** [from Ellen Rony's Domain Name Handbook pages]
**CS Democracy Project** [CPSR ICANN education project]
**Gouverance de l'Internet** [A French site]
**ICANN channel** [A Geman site]
**AfrICANN** [Forum promoting African interests within ICANN]
**Auerbach diary** [This and remaining 3 are MAL Directors pages]

**Campos home**
**MKatoh.net**
**Müller-Maguhn home**

The next Box: Membership Issues with the following links: NAIS Project; Study Committee HomePage; NAIS data request; Bildt committee charter; Members'Forum (ICC); ICANN (non) Members page; ICANN Membership Election site.

Finally boxes devoted to recent links at ICANN Blog for which Brett Faucett deserves much congratulation and to civil liberties issues summarized under the heading: Your Rights on Line.

## Stacked WHOIS Committee Issues First Report

Posted by michael on Tuesday, March 06 @

[**Editor's Note**: The article below is filled with links to other urls. Underlined words indicate links.]

The Whois Committee today issued its first report. [WHOIS, for the uninitiated, is the lookup function by which anyone on the net can see the registration data of any registrant for a domain name in a gTLD. See also better-whois.com.] This interim report is fairly tame, but the Committee has a lot left on its agenda.

A long time ago now - last July - the ICANN staff let slip that it had created a secret advisory committee to advise it on WHOIS issues. This was odd, since ICANN's bylaws require it to be open and transparent. But, not to worry! As ICANN Vice-President, Secretary, and General Counsel Louis Touton explained to Roving Reporter Ted Byfield in October, 2000, it was not a real committee, just "an informal group of people/registrars who have been looking at technical/operational aspects of implementation of the Whois provisions of the ICANN-NSI Registry Agreement and the Registrar Accreditation Agreement."

By December, 2000, the body's cover by now completely blown, Mr. Touton was calling it a committee. He even wrote it a letter asking the "committee" to opine on a large number of questions. ICANN gave it a highly uninformative web page, that doesn't tell you who appointed the members, or when it was formed. It isn't updated much either: On March 6, 2001, the page announced the committee hoped to wrap up its work by January 31, 2001.

The membership of this committee is truly bizarre. Indeed, the word "rigged" comes to mind. Read on...

The WHOIS committee membership, the former "informal group of people/registrars who have been looking at technical/operational aspects of implementation," includes only three groups: Four registrars (CORE, Domain Bank,Melbourne IT, NSI Registrar) The sole current gTLD registry (Verisign) Two leading intellectual property heavyweights and a major telecom:

the Motion Picture Association of America The Recording Industry Association of America Verizon

Depending on what you believe, this group exists either to give the ICANN staff "advice on specific issues that have arisen in the implementation of existing Whois provision for registrars" or to cook up new policies on WHOIS.

If the mission of the committee were purely technical and advisory on implementation issues, the exclusion of the most constituencies and the public might be understandable since this primarily concerns registrars and registries. The inclusion of major intellectual property interests, and of none of the other groups affected by WHOIS, demonstrates that something untoward is afoot. If the objective is to include the perspective of "people" affected by the availability of WHOIS, then membership in this group shouldn't be limited to bodies whose policy agenda is to maximize access to the personal data of every registrant. Other views should be represented too.

In the past, when presented with similar arguments regarding other imbalances in representation, ICANN tends to point to the availability of email-based comment fora. That, quite simply, is bunk. There is an enormous difference between being at the table and setting agendas, being a commentator who either is in the dark or comments after the fact. And if you doubt that, just try establishing an ICANN committee on WHOIS with several consumer representatives but no one from the noted Internet technologists at the Motion Picture Association of America.

Looks like your registrant data is going to remain a lot more open and transparent than ICANN's procedures.

[To respond, click the "Send Your Comment" button in the yellow box to the right.]

## Snookered Again

Posted by jon on Monday, March 12 @ 21:53:06 MST Contributed by michael

A few days ago, ICANN CEO Mike Roberts was explaining that the Board had no set agenda for its Melbourne meeting. People who were getting excited about all the last-minute documents were over-reacting. The usual whining from ICANN-bashers. Well, it's even worse than it seemed: we're stupid, and we've been snookered again.

With no warning to anyone, the ICANN staff pulled a bunch of resolutions out of their pockets at the last minute. There was no public notice. No advance publication. As a result, the entire public comment period the day before the Board meeting was little more than a pointless farce, since no one except the staff (and maybe the Board?) knew what was on the agenda, and almost no one had time to wade through the pile of documents.

The issue of the new gTLD contracts is especially rank. (Recall, as you read all this, that ICANN is supposed to be "open and transparent".)

Let's recall the bidding. The gTLD contracts were supposed to be agreed by Dec. 31, 2000. That would have given us a couple of months to digest and debate them. And make no mistake: these contracts are important if only because they are being touted as "template" agreements that should shape the relationship of future gTLDs to ICANN. Furthermore, these contracts further enshrine the

UDRP (whose review was promised for last years, but has yet to materialize), even as it becomes increasingly clear just how unfair that process works out to be. They undermine the first-come-first serve rule of domain name registration that was always the Internet tradition, and give priority rights to (some) corporations over individuals and other corporations. (You might think such a major change in the rule that has operated since the start of the Internet might be worth some debate?). They further create a new class of reserved names that can't be registered by anyone ? a list negotiated in secret between the ICANN staff and the new gTLD operators without one wiff of public participation.

These contracts were produced in a manner that made informed public participation impossible. And they were negotiated only with the parties that had the most incentive to cave to ICANN. Divide and conquer. Deals in the dark. That's just how ICANN likes it.

Much of these agreements appeared earlier this month; key parts of the agreements in some of the appendixes were published in the last few days. Some parts remain to be published. What we've seen so far, however, seems likely to reshape Internet governance in more fundamental ways than the VeriSign contract. But the VeriSign deal got most of the attention, so the gTLD contracts slipped under the radar.

Mike Roberts deserves a lot of credit here. He as much as promised that the Board wasn't going to decide this issue. You may recall, only a few days ago, Mr. Roberts responded to critics by saying "None, repeat none, of the items [including the contracts] has been posted for action in Melbourne" And, none of them were, even 48 hours before the meeting. Or 24 hours before the meeting. And, in the end, they didn't decide the VeriSign contract, but they decided pretty much everything else. Most importantly, they wrote the staff a blank check on the critical matter of the contracts with the new gTLDs. And ICANN can laugh all the way to the bank, since in addition to giving it substantial regulatory powers over the new gTLDs, they are going to pay ICANN $3.5 million next year, with the possibility of 15% annual increases thereafter. Here's the key resolution that passed with only minimal debate:

Whereas, in resolution 00.89 the Board selected seven proposals to operate or sponsor top-level domains for negotiations toward appropriate agreements between ICANN and the registry operator or sponsoring organization;

Whereas, in resolution 00.90 the Board authorized the President and General Counsel to conduct those negotiations on behalf of ICANN, subject to further Board approval or ratification, to enter into appropriate agreements;

Whereas, the base agreements have been negotiated with the four selected unsponsored top-level domain registry operators (NeuLevel, Afilias, Global Name Registry, and RegistryPro);

Whereas, the base agreements and many of the associated appendices, as completed and agreed by the negotiators, have been posted for public comment;

Whereas, the Board has received a presentation from the General Counsel and the proponents on the progress and results of their negotiations;

Whereas, comments from the public have been received on a web-based public comment forum and at a Public Forum held on 12 March 2001; [This is an especially cruel clause - most of the public comment time was consumed by prepared presentations by invited speakers. There was no notice this decision was coming down the pike. And there's not much indication anyone on the Board actually paid attention to the public comment, and the public couldn't have read all the agreements since they weren't published! In any event, apparently not one iota of these agreements will be altered in light of all this comment.]

Whereas, the Board has considered the posted agreement and appendices, the presentations, and public comments and finds that approval of the agreements is necessary and appropriate to further ICANN's purposes;

It is therefore

RESOLVED [01.__] that the President and General Counsel are authorized and requested to complete negotiation of the remaining unsponsored top-level domain appendices as soon as feasible and to post the resulting appendices on the ICANN web site, along with any minor corrections or adjustments to the base agreement and appendices as already posted;

RESOLVED [01.__] that the Board shall be notified of the complete posting of the agreement and appendices for any of the four unsponsored top-level domains (.biz, .info, .name, and .pro) and after that notification seven days shall be allowed for Board members to make any additional comments to the President and General Counsel;

RESOLVED [01.__] that in the absence of the request of any Board member to the contrary based on policy considerations, the President is authorized to sign the posted agreements after the conclusion of those seven days; and

RESOLVED [01.__] that upon signature of the agreements the President is authorized to take such actions, including causing reports to be made to the United States Department of Commerce, as appropriate to implement the agreements.

You know, it's hard to escape the conclusion that we are stupid. After everything, we believed we'd at least get notice, and a real debate. It seemed obvious that there was no way these contracts could be approved in Melbourne since almost no one -- including, I'd bet, most of the Board -- had time to read them. Indeed, those stalwarts of the IPcommunity Marilyn Cade and Steve Metalitz told the public forum that more time was needed to digest the contracts. No one disputed that. And it seemed obvious that if the Board were going to vote on the contracts there might have been some public notice of that fact.

There is obviously value in rolling out TLDs quickly. But at the expense of laying a decent foundation for how ICANN will regulate TLDs? Is there any way to salvage anything of lasting value from this mess?"

To conclude: the *COOK Report* finds the new ICANN Watch pages to be an exceptionally valuable set of resources that should be used by everyone who cares about the future of the Internet “as a global resource of incalculable value,” and who values having “a central point of reference, a kind of hill overlooking the often-chaotic information landscape, from which anyone seeking a better understanding of these developments can survey the ever-changing terrain.” However, if someone does not soon move against ICANN, we may find that we have lost.

# Optical Border Gateway Protocol Now Internet Draft

**Editor's Note**: On Thursday March 8 we conducted an email interview with Marc Blanchet. Marc is a principal in the Quebec consulting firm Viagenie and has done much of the implementation work on OBGPfor Canarie and Bill St Arnaud. In early March Marc and his Viagenie colleague Florent Parent, and Bill St. Arnaud submitted to the IETF a formal Internet Draft <http://www.ietf.org/internet-drafts/draft-parent-obgp-01.txt> for consideration at the spring IETF in Minneapolis.

***COOK Report***: What do you expect or hope will happen at the IETF meeting in Minneapolis?

**Blanchet**: we will be presenting our draft. This will make people aware of our work. We think it is focused on the interests of the working group work. We would like to see our work as being included in the working group work. We are also looking for feedback on the proposal.

***COOK Report***: The name and purpose of the working group that you are presenting to is what?

**Blanchet**: IPO: IP over Optical. See http://www.ietf.org/html.charters/ipo-charter.html

***COOK Report***: What may happen in IETF AFTER the Minnesota meeting?

**Blanchet**: Having gotten more feedback from people, we expect to undertake revisions of the specification in order to have the document accepted as a working group document.

***COOK Report***: Over what period of time? 3 to 6 months?

**Blanchet**: I have no idea. As I did explain, it often happens that the draft goes live for a "long" period of time before freezing as an RFC. It is right now too soon to know how much time we will spend before having a ready-to-publish spec.

***COOK Report***: Would you describe some of the responses to the Internet draft. Just a general description is fine.

**Blanchet**: People from different vendors seem interested. We will actually meet some of them in Minneapolis.

***COOK Report***: What kinds of products may have the code implemented in them?

**Blanchet**: We don't have information about the internal strategic plans of the vendors.

***COOK Report***: I mean very general. What kind of devices can use it? Switches? Routers? Anything else?

**Blanchet**: the way the industry is doing things regarding optical and IP, it looks like we will probably need to invent a new word for describing these devices.

***COOK Report***: How long before the first commercial implementation may be expected?

**Blanchet**: I can't guess. It is known that both CA*Net and Internet2 will be probably issuing a Request for Proposal (RFP) soon, so there will be market demand. In that context, this generally means vendors working on an implementation.

***COOK Report***: Would you explain in one or two sentences why this will cause demand? (Please forgive my minimal knowledge.)

**Blanchet**: If today there is no buyer for a product, then there is still chance that vendors will do the product and see what happens. (The internet has this kind of history: the web itself, napster,) But if today there are important buyers for a product, then there is more chance that vendors will do a product.

I was saying that the second is happening, so this probably means that there will be some implementation soon available! But as said, I cannot disclose vendors internal plans.

***COOK Report***: Do you have a target date for the proof of concept code for testing in the Quebec research network switch?

**Blanchet**: During the coming summer.

## Network Working Group Internet Draft

Expiration: August 2001
Florent Parent,
Marc Blanchet, Viagenie
Bill St-Arnaud, Canarie
March 2001

### Optical BGP (OBGP): InterAS lightpath provisioning draft-parent-obgp-01.txt

#### Abstract

This work investigates inter-AS lightpath provisioning using BGP. BGP is currently deployed across the different autonomous systems on the Internet, so investigating how a BGP approach to provision lightpaths across autonomous systems is of great interest. This work proposes extensions to BGP to this end.

#### 1. Introduction

Much of the current work in IP optical focuses on using interior gateway protocols such as ISIS and OSPF with TE extensions for routing and GMPLS for signaling [ISIS-TE], [OSPF-TE], [CRLDP], [RSVP-TE].

This draft considers optical lightpath provisioning at the inter-AS scope. Other protocols may be use inside the autonomous system to control the actual optical cross-connect (OXC).

BGP [BGP] is currently deployed across the different autonomous systems on the Internet, so investigating how a BGP approach to provision lightpaths across autonomous systems is of great interest.

The goal of this work is to propose a BGP approach to lightpath provisioning.

**2. Scope**

The traditional networks are carrier based and the carrier manages the customer transit connectivity. Customers are now acquiring fibre, optical switch ports and wavelengths and this changes the traditional model of network peering.

Customers in sites can now own their own wavelengths, and eventually optical switch ports. Providers can deploy an optical infrastructure and sell wavelengths and optical switch ports to their customers. Customers now have multiple wavelengths and optical ports from one or many providers.

This ownership of wavelengths and optical ports now brings new operational requirements where customers at the edge need to control a subset of lightpaths within another network's wavelength cloud (provider) so that they can manage their own lightpath routing within that cloud.

This new model allows distributed Internet Exchange facilities using the exchange and trading of lightpaths between networks to minimize the need for hierarchical network architectures to interconnect peering networks. As an example, CA*net4 optical network in Canada could provide lightpath transit to Renater/France and Wide/Japan networks.

The scenario that is considered in this document is where autonomous sites own their wavelengths and optical switch ports. End customer have "virtual" control over its optical switches/ports/wavelengths.

The document covers inter-AS provisioning using BGP. Inside an autonomous system, other protocols can be used, such as OSPF or ISIS with TE extensions for routing and GMPLS for signaling.

**3. Protocol operation**

It is assumed that a BGP peering is already established between participating sites, either using a non-optical path or a pre-configured optical path between sites. This BGP peering will be used in the following description.

The (egress) OXC are BGP speakers. The OXC and the BGP router are closely tied together in the sense that information received from BGPwill be used to establish optical cross-connects inside the OXC.

The sites are eBGPpeers. This document doesn't specify any intra-AS routing and signaling protocols for lightpath provisioning, although interaction between the inter-AS and intra-AS protocols will need to be defined.

The protocol proposes two phases.

The first phase is the lightpath reachability phase. During this phase, sites advertise through BGP the availability of the optical lightpath to their site. These announcements will contain information on the OXC and the available lightpath through that OXC. The information will be encoded using multi-protocol BGP extensions and extended community.

This first phase will allow sites to build up a "lightpath RIB" that will be used to determine if a lightpath is feasible across a number of OXC in different sites.

The second phase is the lightpath establishment. This phase uses the information received from the lightpath reachability phase and then uses a BGP update message to communicate the lightpath establishment to the OXC sites on the path. The information will be encoded using multi-protocol BGPextensions and extended community.

4. Encoding Optical Lightpath information in BGP

The BGP Multi-protocol extensions [BGP-MP] allow BGP to carry routes from multiple "address families".

This document proposes to use MP-BGP extensions to encode the NLRI such that the necessary optical and routing information can be propagated in BGP. Figure 1 presents the MP_REACH_NLRI attribute format.

**Editor's Note**: for remainder of draft see: http://www.ietf.org/internet-drafts/draft-parent-obgp-01.txt

# Letter to the Editor

# DANTE Objects to Description of Them by Kees Neggers

On March 1 we received the following:
From: Dai Davies
<Dai.Davies@dante.org.uk>

Dear Mr Cook

You recently published a report in relation to the Dutch Research Network, SURFnet, where you printed an interview with Mr Kees Neggers, their Managing Director. In this Mr Neggers makes a number of statements in respect of the GEANT project, which are incorrect.

In particular:-

1. Mr Neggers states that the entire Consortium of 26 National Networking entities in Europe has to agree on everything in the context of GEANT. This is not the case. A simple majority is required for most decisions and major commercial decisions require a two thirds majority. The number of votes that each of the partners has relates directly to their capacity commitment to the network.

Therefore, countries with high speed accesses have significantly more say in the overall direction of the project.

2. Mr Neggers doubts that GEANT is state of the art. The whole GEANT initiative is, indeed, aimed at being at the state of the art. It has, as part of the project, an important development programme which covers, initially, the implementation of services such as IPv6, differentiated Quality of Service etc. and commitment to this ongoing development activity is an important element of the overall GEANT network.

3. With respect to Dutch connectivity, it is not apparent that SURFnet will simply be able to connect to GEANT. As they are not a member of the consortium they are taking no part in the commercial and technical risk associated with the activity and will require to reach agreement with the consortium before any service can be provided in the Netherlands.

The Netherlands is a small, homogeneous, country and it is relatively unchallenging to make technical progress in such an environment. It is much more interesting to implement state of the art technology in a federated environment with thirty networks and a pan-European interconnection. This is a much truer reflection of the challenges of implementing and developing advanced Internet services.

Best wishes, Dai Davies, General Manager, DANTE

On March 1, 2001 ***COOK Report replied***: Would you like us to publish this as a Letter to the Editor in about 2 weeks?

**Davies**: Yes this would at least set some of the record straight. It would also be nice to have the opportunity to give your readers a rather more accurate and positive view of the GEANT project, than that which they might derive from your interview with Kees Neggers.

***COOK Report***: giving Kees Neggers an opportunity to reply?

**Davies**: He is welcome to try and reply to these points, if he has anything to say.

***COOK Report***: On March 1, Kees replied: "I see no need to respond at all."

Watching what SURFnet did was a pretty good response to Davies. It established the fastest external connectivity of any research network in the world. On March 6 SURFnet Press Releases were posted at http://www.gigaport.nl/en/en_main_act.html

## Teleglobe Preferred Supplier for Global Internet Connectivity GigaPort and Lambda-connection StarLight

"Utrecht, 5 March 2001 - In a public tender GigaPort has awarded Teleglobe 'preferred supplier' status for supplying a number of services for the global Internet connectivity of the GigaPort network. With Teleglobe GigaPort is now drafting a contract for the period of July 2001until the end of 2002."

"In a public tender in compliance with European regulations GigaPort received thirteen offers with proposals for the realization of Global Connectivity. Teleglobe was made preferred supplier for part of this tender. The aim of the cooperation between GigaPort and Teleglobe is the realization of global connectivity from the SURFnet5 backbone location at Hempoint in Amsterdam with a speed of 1 gigabit per second. GigaPort will also realize a 2.5 Gbit/s lambda connection with Teleglobe between SARA(the other SURFnet5 backbone location) and the StarLight (successor of STAR TAP) in Chicago, to be able to experiment with new types of technology for a completely optical Internet. This connection can be upgraded to a 10 Gbit/s capacity in the course of 2002."

"Finally, GigaPort and Teleglobe will realize a Virtual Private Network connection from Hempoint to the SURFnet Point of Presence in Chicago, which will be built within the GigaPort project. The PoP in Chicago will replace the current SURFnet PoP in New York City, and like the current PoPit will be connected to the STAR TAP at a speed of 622 Mbit/s. Through STAR TAP GigaPort has connections to several high-speed networks in North America and Asia."

Finally on March 9, **Bill St Arnaud** posted to the Canarie list:

"SURFnet - the Netherlands research and education network recently announced their new SURFnet5 network. We congratulate our colleagues at SURFnet for their vision and world leadership in this area. We especially congratulate the Dutch government for the strong financial support they have provided to the GigaPort and SURFnet5 project which has now allowed Netherlands to become a world leader in next generation optical Internet research and education networks."

"SURFnet5 will have 32 lambdas at 10 Gbit/s to start with. In addition SURFnet5 has selected Teleglobe as its preferred supplier for dedicated wavelengths across the Atlantic from the Netherlands to the new StarLight optical Internet exchange in Chicago. This will be another world's first and will significantly position the Netherlands to be a global hub for next generation bandwidth intensive applications."

"The core of SURFnet5 consists of two backbone locations in Amsterdam. Each backbone location consists of two Cisco 12416 routers. 15 concentrator locations will be connected to this core. Each concentrator location will have one 10G connection to each of the backbone locations. Each lambda at 10G will be an unprotected lambda inside the DWDM network of Telfort, the carrier supplying the transmission facilities for SURFnet5. Resilience will be introduced on the IP-layer."

For more details please see www.surfnet.nl or www.gigaport.nl

# Executive Summary:

## LayerOne, pp. 1 -14

We interview Alexander Muse, President and CEO of LayerOne which is offering optical interconnects for carriers at physical layer of the OSI reference model. LayerOne, located inside carrier hotels, uses the Ciena Core Director switch as a service to enable many different carriers to connect their fiber strands and provision circuits for each other.

LayerOne takes between 2,000 and 15,000 square feet of floor space insides a carrier hotel. They then bring in between four and 30 thousand strands of fiber into what they call a Nexus Bandwidth Exchange. The fiber is connected to a complex array of cabinets. To enable this they have a framework of ADC FEC boxes and the ADC FL2000 SC Connectors ELF bay. For electrical connections at the DS1 and DS3 level they use a customized ADC Entraprise Frame. Optical to optical connections at levels ranging from dark fiber to OC-192 they make via ADC's Next Generation Fiber Frame. Ethernet connections are handled from a router/switch.

From this equipment fiber is tied into the Core Director. Instead of having customers bring in bays full of expensive Sonet equipment LayerOne asks for the delivery of fiber strands in 96 or 192 strand cables. Initially they will connect two of them to the Ciena Core Director - each of which can handle up to 256 OC48s and each of which uses transponders and tunable lasers to do its work. Within the Core Director those would be then lit to a level of OC 48 or OC-192, depending on how their customers would want to interface. From there LayerOne grooms fat customer pipes to electrical STS -1s. It then maps those to the other providers to whom they want to connect. LayerOne charges them a flat monthly fee for connecting to it Core Director. The size of the fee depends on the bandwidth of the connection which can range from DS-1 to OC192.

But what is this grooming all about? As light reading explains, <http://www.lightreading.com/document.asp?doc_id=4118> the Core Director is capable of grooming. That means it can set up any size of pipe across a network by combining any number of small STS1 (51.8Mbit/s) connections, ranging in size from 1 to 192." Grooming "slashes the number of boxes that service providers need to buy and maintain; it also helps them provision services faster and use bandwidth efficiently."

LayerOne is open in five locations and plans 15 or more by years' end. In the Bryant Street Dallas colo it is currently interconnecting 30 carriers. These range from Qwest and Level3, to SBC and Verizon, to Yipes and Telseon, to MFN and lesser knowns like Global. Metro.

Finally we conclude with a three page commentary on LayerOne by New York City consultant Frank Coluccio who paints a contextual picture of the LayerOne developments and who, emphasizing a point made by Muse, notes that "the capabilities that Layer One demonstrated were possible, were not capabilities that their vendor, Ciena, would have initially guaranteed, much less expect LayerOne to attempt to accomplish."

## Net Access, pp. 15 - 25

We interview Avi Freedman who relates his experience during the past year as he has sought to upgrade Net Access metro infrastructure from carrier provided to self provisioned using dark fiber. The network that he founded in Philadelphia in 1992 is now a 10 million dollar a year business supplying customer circuits from Boston to New York City, to Philadelphia to Baltimore to Washington DC.

Freedman's problem was to determine whether it made sense for him to buy Gigabit Ethernet from Telseon or Yipes!. Freedman says: “Gigabit Ethernet is something that we've actually found that people don't want primarily because it is difficult to be certain one is not getting shortchanged.” "The fact is if I give you an OC12, you know you are getting an OC12 and there's nothing I can do to short change you on the matter of how many bits you can pump through it. You know that you have your full circuit even if it is being time sliced and sent over my WDM infrastructure. But if I give you a Gig E, you have no clue whether that's a VLAN on a big trunk or whether it is a real Gig E dedicated all to me, because the technology allows you to sell 10 people a GigE, put a 10 meg Ethernet in the middle and then sell you all Gig Es at the other end."

Freedman has found that it is most cost effective for him to lease dark fiber from MFN than to buy Gigabit Ethernet service from Telseon or Yipes! because they won't let him put his own measurement tools on the circuits. Under such circumstances the cost of their bandwidth is close to what it would cost him to lease dark fiber to supply an equivalent amount. To be competitive he figures the cost should be about one half of what he could provide it for himself. He also looks at Cogent which he believes to be selling to ISPs at rates that are below its cost. He questions how long they can stay in business.

He gives a great deal of useful pricing information and talks about scaling issues involved in the cost of providing the largest circuits (OC 192s). Provisioning OC48s is now more cost effective. As he puts it: "If you can do 32 OC48s and get another fiber pair, do 32 more OC48s, it can be cheaper (especially if you're trying to make your capital last) to just do this." He also explains that his customers have told him that they also want plain OC circuits without IP on them because what the carriers used to provide in seven to 30 days five years ago now takes six months or longer.

The article concludes with an exchange between Telseon Vice President Bob Klessig and Avi Freedman. A high light --**Klessig:** “It is safe to assume that the traffic between the routers is highly concentrated and that the routers will be in place for a long time. As Mr. Freedman says, this is a great scenario for a dedicated link such as dark fiber if it is available. But Telseon is offering a switched service. It provides for new data connections in minutes to hours between any two points on a Telseon metropolitan network. Bandwidth choices are highly granular and quickly changed.”

**Freedman:** “There are no applications or software demanding highly granular and changeable bandwidth today that I know of from Net Access or Akamai. The base cost of Telseon and Yipes! 1mb-over-100mb/1000mb services are as much as customers will shortly be able to pay Net Access, and probably Enron and others for a 100mb provably dedicated pipe. Analyses and promises are no substitute for proof of infrastructure where the possibility for aggregation follows.”

## ICANN - Verisign, pp. 26-35

Those who established ICANN got much of their early support from pledging to reign in what they painted as a very nasty Network Solutions monopoly in the dot com and other gTLDs. Now operating under the guise that the Internet industry can be trusted to regulate itself, ICANN and Networks Solutions successor, Verisign started in secret last summer negotiations which on March 1, 2001 effectively reimposed the monopoly. While we offer our own analysis we also recommend that of Brock Meeks: Dot Com Hocus Pocus -The Remaking of a Monopoly http://www.msnbc.com/news/540693.asp.

The Australian Board meeting ended with the Board gramting ICANN staff a blank check to make decisions for which they the Board should take responsibility. As Michael Froomkin wrote: Well, it's even worse than it seemed: we're stupid, and we've been snook-

ered again. With no warning to anyone, the ICANN staff pulled a bunch of resolutions out of their pockets at the last minute. There was no public notice. No advance publication. As a result, the entire public comment period the day before the Board meeting was little more than a pointless farce, since no one except the staff (and maybe the Board?) knew what was on the agenda, and almost no one had time to wade through the pile of documents.

Meanwhile Vint Cerf as Board chair acted to involve himself as Versign's interests. Verisign CEO Sclavos wrote to Cerf on Feb. 28: "We also appreciate your commitment to seek formal Board approval for an appropriate extension of the time under the existing agreement should compliance with Section 23 be necessary. But we are hopeful that by working with you, and the Internet community, including members of the ICANN Board, we will all see these new agreements approved and successfully implemented."

A BWG attorney commented: Section 23 is the one that says that the com/net/org registry agreement expires in November 2003 unless Verisign divests the registrar by May 2001. Read the third sentence of the above paragraph. Sclavos is saying that Cerf has made a personal commitment to him that, if the ICANN Board does *not* approve the proposed agreements, Cerf will go to the Board to get the May 2001 divestiture deadline extended.

We also include a review of the completely revised ICANN Watch website <http://www.icannwatch.org> This is now by far the best site on the web for tracking what this renegade Internet regulator is doing.

## OBGP Internet Draft, pp. 36-37

OBGP is now an Internet draft <http://www.ietf.org/internet-drafts/draft-parent-obgp-01.txt> for consideration at the spring IETF in Minneapolis. We briefly interview one of the draft authors Marc Blanchet about what to expect in coming months.

## DANTE Geant Letter to Editor, pp. 37-38

Dai Davies General Manager of Dante suggests SURFnet criticism was unfounded. Kees neggers informed us he saw no need to respond. On March 6 SURFnet established the fastest external connectivity of any research network in the world. SURFnet Press Releases were posted at http://www.gigaport.nl/en/en_main_act.html



## Subscription Rates

Choice of either ascii or Adobe Acrobat (PDF) format 1. Individual; College or University Department; or Library; or Small Corporation - $295 2. Corporate - (revenues $10 to 200 million a year) - $375 3. Large Corporate- Revenues of $200 million to $2 billion per year - $475 4. Very Large Corporate- Revenues of more than $2 billion per year - $575

Site License: The right to distribute ascii and PDF via email to all employees of corporation. 5. Small corporate: $450 6. Corporate: $700 7. Large Corporate: $950 8. Very Large Corporate: $1200 . 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, USA**