



The COOK Report on Internet



© Cook Network Consultants

AT&T Explains Its Optical Network Architecture In Search of Business Models for a Slower Growth Era -- Carriers Try to Cope with Hang Over of Fiber and Debt after the Great Optical Build out of the 90's

Editor's Introduction

In this issue we look at some of the bigger picture consequences of the growth of the internet and the web. The look is necessary if decision makers are going to understand some of the possible outcomes inherent in the economic downturn of the current market. In the mid 1990s there were only three nationwide fiber optic based telecommunications transport networks in the United States (AT&T, MCI and Sprint). Prior to the emergence of the Internet, it was thought that the fiber then in the ground would take many years to fill to capacity. We can remember the forum on access to a government sponsored National Research and Education Network (NREN) we lead at the US Congress Office of Technology Assessment on December 11, 1990. Steve Wolff in his capacity Director of the NSFnet assured the meeting that no new fiber would be needed to support the NREN since available fiber was far from being filled. Only five years later the picture had radically changed.

As AT&T researcher John Strand points out, by 1995, the optical transport networks of AT&T, Sprint, and MCI accounted for 75% of total inter city fiber deployment in the US. By the year 2000, MCI, Sprint and AT&T fiber had shrunk to less than 1/3 of the national total with the remainder belonging to no less than 39 new national carriers. And, as the year 2000 began, there was an estimated total of 400,000 route miles of

fiber with an average 46 strands of fiber per cable.

By 1995 it was obvious to everyone that the Internet was revolutionizing telecommunications in general and was responsible for huge increases in data traffic in particular. With the impact of the World Wide Web hitting the net like a tidal wave in 1995 and the commodity backbone startups on April 1, 1995 as the NSFnet backbone closed down, it became impossible to get accurate statistics on bandwidth usage. However, best estimates pegged it as exploding. It was doubling every 3 to 4 months. Over the next several years, UUNET's Mike O'Dell's continual theme in speeches at conferences was continued doubling of UUNET's backbone capacity every three to four months.

Riding a Bubble Buoyed by a Bandwidth Myth

Not wanting to be left behind, everyone began to make huge investment in Internet infrastructure. As AT&T researcher Andrew Odlyzko writes "After all, if the Internet grows by a factor of 16 each year, then the first mover may have an unbeatable advantage." Investment flowed into the development of optical network technology and resulted in the technology being pushed by 1996 and 1997 far more rapidly than people suspected would be possible. By 2000 not only had the amount of fiber based route

Volume X, No. 5, August 2001
ISSN 1071 - 6327

miles grown by nearly 10 fold in about five years, but the amount of data that could be sent over each fiber had increased by 100 fold with promises of potential 1000 fold increases.

A gold rush mentality developed and money poured in with every new player seeking to capture a first mover advantage in what people sensed was a fundamental restructuring of telecommunications. Suddenly people noticed that all the old assumptions for planning and investment no longer held as the patterns for data traffic and structures needed to cost effectively scale such traffic were not necessarily compatible with the investment in the kind of technology needed to grow the vastly differ-

On the Inside Issues in Network Access

AT&T Net Design	pp. 1-11
Level 3 Intro	pp. 12-17
Peering	pp. 18-21
DDOS attack	p. 21
ICANN	pp. 22-27, 30
Executive Summary	pp. 28-30

ent patterns and needs of voice traffic. To make matters worse for planners and for allocation of financial resources, while data traffic was growing at least ten times as fast as voice traffic, voice traffic could be sold for seven times more per megabyte than data traffic.

There was only one problem. The bandwidth growth estimates were wrong. Odlyzko and his colleague Kerrey Coffman first questioned the growth estimates in the fall of 1998. By the winter of 2000 – 2001 they had done some serious studies. Readers should look at the short paper from which these quotes are taken: Internet growth: Myth and reality, use and abuse, Andrew Odlyzko, AT&T Labs - Research <http://www.research.att.com/~amo/doc/recent.html>. Readers should also note that Andrew however pointed out to the COOK Report that the “myth and reality” piece is just a short blurb based on a much more detailed and scholarly work, namely the paper “Internet growth: Is there a “Moore’s Law” for data traffic?,” (joint with Kerry Coffman), which is available at <http://www.research.att.com/~amo/doc/networks.html>.

According to Odlyzko: “Almost all references to Internet growth claim astronomical rates of increase; the usual phrase is that “Internet traffic is doubling every three months.” Even serious observers echo such claims. For example, former Federal Communications Commission Chairman Reed Hundt recently wrote, “In 1999, data traffic was doubling every 90 days. . . .”

“Amazingly enough for a claim that is so dramatic and quoted so widely, there have been no hard data to substantiate it. Indeed, careful scrutiny of existing evidence on traffic by a number of experts suggests that the truth is considerably more modest. Internet backbone traffic in the United States has been about doubling annually for the last four years, and currently appears to continue growing at about that rate. (Doubling is used here in a loose sense to cover growth rates between 70 percent and 150 percent per year.) Although this is extremely rapid growth, much faster than in any other

communications service, it is not anywhere close to the 700 percent to 1,500 percent annual growth rates that a doubling of traffic each three or four months would imply.”

“As often happens, there is a grain of truth behind the claims of Internet traffic doubling every three or four months. Such growth rates did prevail for a short period during 1995 and 1996. This brief period of extraordinary growth seems to have colored popular perceptions. The belief that Internet traffic could continue “doubling every three months” all this time shows an astonishing degree of innumeracy, the lack of simple quantitative reasoning. At this rate, traffic would be increasing by a factor of 16 per year. Hence, from the end of 1994 to the end of 2000, it would have grown by a factor of almost 17 million. “

“The Internet backbone at the end of 1994 was somewhat larger than two T3 (45 megabit per second) links crossing the continent. Growth by a factor of 17 million would have produced a network with over 600,000 OC48 (2.5 gigabit per second) links from coast to coast, far more than all the existing fiber strands could carry. “

“The “doubling every three months” over the last six years story is simply not consistent with reality.” “The myth of Internet traffic that doubles every three or four months is dangerous. It leads to bad decisions. It surely helped inflate the current bubble in optical networking stocks. After all, if demand is outpacing supply of transport capacity, then money making opportunities are virtually limitless,” Odlyzko concluded. We pointed out Odlyzko’s work to Avi Freedman who has planning responsibility for Akamai’s global content network. He agreed with Odlyzko’s conclusions, saying that “anecdotal evidence is that things grew about 100% in each of the last two years, but no faster, and may be growing more slowly now (though not declining). I base these conclusions on looking at how much of an ISP’s traffic Akamai is, and how much traffic we do, and how much traffic each ISP sends as a whole. Also, since that doesn’t catch Napster and other non-http and non-streaming proto-

cols, I looked at traffic from about ten ISPs who’ve remained steady in the size of their user base.”

By 2000 an average growth of about 100% a year meant that in some instances available capacity had grown to exceed actual usage substantially. Unfortunately by the time many people began to be aware of the mistaken assumptions about bandwidth growth, the global telecom industry was coming of a five year long binge of spending. With the end of the investment boom of 1995 to 2000 in the winter of 2000 - 2001, the telecom industry found itself saturated in debt. It’s debt position was well chronicled by Ravi Suria in private papers for banking and investment clients at the end of 2000 and in a lengthy public interview in early 2001. As we have previously reported, Suria argued persuasively that the debt positions of the majority of telecom players are unsustainable and that many bankruptcies would follow. This then is the context in which, those who have invested in new fiber infrastructure, must try to figure out what they should charge for bandwidth derived from that infrastructure. It is also the context in which consumers of bandwidth must try to decide how much to pay now and, even more important, how to allocate capital to pay for their bandwidth needs over time.

First Questions of Optical Network Design and Interoperability

The problem is that as we enter the summer of 2001, the ability to make sound judgments in the acquisition of bandwidth over the long term has vanished. A decade ago someone contemplating a 10 year IRU on a circuit of specified bandwidth could make a reasonable judgment of what he should pay in year five and year ten. Today it is difficult to make a sound judgment about bandwidth a year in advance. Without that ability, some bandwidth-based businesses are simply too risky to launch. Consequently, under such circumstances it is more difficult for the carriers and the new green field telecom players to recoup their investments in their infrastructure. Some service

providers are even abandoning facilities-based strategies all-together and are turning to "carrier's carriers" for bandwidth.

In this new era of tightened capital investments, it has become obvious to many players that those who can integrate their new investment in fiber and optronics into their infrastructure most cost effectively would be best equipped for survival and growth. With the fiber infrastructure in place and the optronics to power the 'bits' becoming mature, the technology has made new market opportunities possible. Consequently, the race is now on to take advantage of the new technology in order to create optical transport networks that are cost-effectively meshed with the growth of data traffic and the many new markets of high bandwidth services that the Internet has made possible.

A key element to leveraging infrastructure investment is flexibility. Facilities-based service providers must be able to quickly connect new customers and change the topology of the network in response to new services and changing business conditions. This capability, commonly referred to a "rapid provisioning", enables service providers to provide access to all of their installed fiber capacity on an as-needed basis. Part of the problem is that optical networking equipment was originally developed beginning in 1998 on a custom basis for specific customers. Under such circumstances the creation of standards to allow for rapid provisioning to occur across optical networking equipment from different vendors was not a high priority. Likewise, standards to allow data networking equipment cannot talk directly to optical equipment to request bandwidth do not exist. As a result, for the owners of the new fiber networks, full flexibility in leveraging infrastructure is not easily attainable.

But down in the trenches some people knew that they better begin to address these problems. According to <http://www.oiforum.com/> "On April 20, 1998 Cisco Systems and Ciena Corporation announced an industry-wide initiative to create the Optical Internetworking Forum (OIF), an open forum focused on

accelerating the deployment of optical internetworks. The founding members of the forum were AT&T, Bellcore (now Telcordia Technologies), Ciena Corporation, Cisco Systems, Hewlett-Packard, Qwest, Sprint and WorldCom (now MCI WorldCom). The mission of the OIF would be to provide a venue for equipment manufacturers, users and service providers to work together to resolve issues and develop key specifications to ensure the interoperability of optical networks." Building an administrative and legal framework for the forum took the remainder of 1998. "By January 1999 the following technical working groups were up and running with official charters and chairs: Architecture, Physical and Link Layer, OAM&P (Operations Administration, Maintenance and Provisioning)." By March of 2000 the OIF had 169 member companies. And by early 2001 a carrier working group and a signaling working group had been added to the initial three working groups.

In parallel with the OIF, more formal standards bodies like the IETF and the ITU have been making important contributions. The IETF's Generalized MPLS (GMPLS) work aims at extending MPLS's provisioning capabilities to optical, SONET/SDH, and potentially other technologies, while the ITU has taken the lead in standardizing the wavelength grid and has proposals for standardizing the structure of the optical channel.

John Strand is Chair of the OIF working group that developed the following definition of the goals that were being sought

"Optical networking permits carriers to provide new types of network services not available with other technologies, enabling sophisticated transport applications of (D)WDM based networks (featuring a variety of topologies such as point-to-point, ring and mesh). These new generation networks provide means for the improved use of network resources and the support of high-bandwidth services. Dynamic bandwidth allocation, fast restoration techniques and flow-through provisioning give birth to an assortment of services."

"Intelligent OTN's contain distributed

management capability and subsume many provisioning and data basing functions currently performed by carrier Operations Systems (OS). This allows the rapid establishment and reconfiguration of connections, potentially reducing provisioning times from months to seconds, thus lowering operating costs and providing the means to set and guarantee SLAs and QoS configured on a per-connection basis to better meet customer's specific needs. "

"The large capacity and great flexibility of such networks enables the support of several degrees of transparency to user traffic at lower cost to the end customer. The new services expected to be enabled as a minimum are bandwidth on demand, point and click provisioning of optical connections, and optical virtual private networks."

"The standardized interface between the optical layer and the higher layer data service layers such as IP, ATM, SONET/SDH enables the end-to-end internetworking of the optical channels for conveying user information of varying formats. The use of standardized protocols will make the benefits of the intelligent OTN's available end-to-end, even if several networks are involved."

In the absence of the necessary inter-operability standards these new optical network services would have to be developed first within the networks of carriers and their greenfield competitors. John Strand and Tom Afferton have consented to talk with us in this interview about AT&T's approach to building a flexible, cost-effective optical transport infrastructure for its own network and how interoperability standards contribute to this process.

John Strand is a consultant for the Optical Networks Research Department at AT&T. He is the chair of the Optical Internetworking Forum's Carrier Working Group and active in the IETF's IP over optics (IPO) Working Group. He also gives many short courses for the Optical Society of America and individual companies, primarily start-ups. He also is a consultant and member of the Technical Advisory Boards of several start-ups,

and will be a visiting scholar at the Univ. of California-Berkeley in 2001-2002. Previously he held various technical and management positions at AT&T in network architecture, network planning, services planning, and software systems engineering and development groups. He received a Ph.D. in mathematics from the U. of California-Berkeley and an A.B. in economics from Harvard College.

Tom Afferton is currently a district manager in AT&T Network Services responsible for Advanced Transport Technology and Architecture Planning. In this role, he establishes the architectural direction for AT&T's domestic metro and intercity transport networks. He also manages AT&T's Advanced Transport Technology Incubator, a collection of resources dedicated to collaborative trials with emerging vendors of innovative transport and optical networking technologies. He received a B.S. in Electrical Engineering from the University of Virginia and an M.S. in Electrical Engineering from Stanford University. He is currently a member of the board of directors of the Optical Internetworking Forum and serves on steering committee of the Optical Fiber Communications Conference.

Dave Johnson, Media Relations Director of AT&T Network Services also participated in the interview.

COOK Report: So how would you describe the process of your work on architecture design?

Afferton: I am responsible for advanced transport technology and architecture planning. I am at the front end of our technology introduction process for the transport network. John is located at AT&T Labs and supports my organization with research looking further ahead in the space and focusing on how it relates to our ongoing services.

Strand: I'm here to give Tom ideas. We have a constant dialogue going about ideas and trends. He does the first-round work in boiling them down into something real.

Afferton: I am responsible for our plan of record for the transport architecture

and other folks in our organization would be involved with the specific network requirements, and the specific network elements. John's department in research works with me in designing the architecture and in looking at new technologies. This enables other folks to follow through on the implementation.

The Carrier Context of Network Design

COOK Report: How would you describe the context in which you are working?

Strand: In the Optical Internetworking Forum group that I chair, the question of how to handle the very different characteristics of voice and data traffic effectively continually comes up. In many ways it's a paradox. We need to support increasing amounts of data traffic at the same time we are dependent on the much larger amounts of income that we derive from decreasing amounts of voice traffic.

Afferton: That leaves us in a situation where we are building a network to deliver services to our customers and where we are anticipating customer needs and are putting in technology in advance of those needs but at the same time are also reacting to the demands of the current marketplace. Our customer needs are very much factored into our network design. It is an iterative process and not just a technology push.

COOK Report: How then would you summarize the environment in which you have to plan the further development of your network?

Afferton: Certainly there is a problem statement that we can start with in looking at our architecture. It is one that is derived from a set of business conditions. We are certainly seeing dramatic growth of the demands on our network. I think that Dave Johnson has some statistics as to how much volume we handle in a given day.

Johnson: On a typical business day the AT&T network handles in excess of 995 trillion bytes or terabytes of data. This is equivalent to the entire printed contents

of the Library of Congress every 30 minutes. In addition to the 995 terabytes of data our network handles 300 million voice calls daily. This is a typical business day and not an especially busy one.

Afferton: We're seeing a very high volume of traffic as we speak. This is happening at the very time in which we're also seeing rapid growth in that traffic. Our growth in voice has become more stabilized. It comes in single digit percentages. Our data services are growing much faster.

Johnson: ATM and frame relay are growing about a 60 percent per year; IP is growing at a rate of about 200% year over year.

Afferton: I think John has some statistics from the labs on Internet growth of about 100% a year after year.

Strand: I talked to Andrew Odlyzko and Kerry Kaufman who confirmed when they look at their data on Internet growth through the end of 2000 they see no slackening in its pace of 100%-per-year growth.

Afferton: You have then rapid growth most of which is coming from the data network but at the same time you have the existing voice services from which you do derive a fair amount of revenue as well. The challenge that comes to me is to evolve the transport network to support all of AT&T's services in a cost-effective way. Since there are a lot of other networks out there that are putting out bandwidth, we need to be competitive with them from a unit cost stand point. In fact, our goal is to be the leader in terms of best-in-class unit costs. And to do this, we look across not just to the long haul and core backbone but also metro and access networks. After all, our total business picture depends on unit cost from all these areas.

Problems of Voice Versus Data

Strand: In trying to accomplish the above goals, we find that Internet traffic characteristics present their own dilem-

ma. Due to the basic nature of packet traffic it makes sense to carry it in very large pipes connecting routers together. Now in the voice world we deal with what we call service providing facilities which are the actual connections between the boxes that are providing the services. In the voice network's case those were 64 kilobit services or, coming out of our 4ESS switches, 1.5 megabit T-1s. This creates the need for an elaborate network to bundle up and manage T1s and T3s which are smaller than many of the data pipes we are now comfortable with.

The key point that I'm trying to make here is to show some of the things that are forcing big changes inside the legacy voice world. Traffic in the voice world appears in relatively modest quantities. We had to do an enormous amount of work to build up voice pipes that would be large enough to be transportable on a fiber network. While in the IP world traffic is increasingly appearing as OC-3s, OC-12 s, and OC 48 s. As a result, this IP traffic bypasses a lot of the infrastructure that was built up to handle voice traffic.

COOK Report: So you've got two different kinds of traffic with, to some extent, two different kinds of infrastructure and therefore part of the problem is how to make the handling both of them cost effectively in a way that is compatible with your business needs and resources at hand. With the one growing and the other a flattening, what you do with them? You have to deal with both. But what is the most effective way of dealing with voice is likely not the most effective way of dealing with data.

Afferton: I think you touched on a point with which I would resonate. Given that the growth is in the data network, for the most part when we're talking about scalability of the network, we're talking about scaling for the growth. We certainly do need to continue to support the voice services and to have an infrastructure that supports such. But we also have to realize that, to a large degree, we have that infrastructure already in place.

COOK Report: And will that infrastructure that is in place work over the optical network OK? It is SONET and all the

things that that implies, is it not?

Afferton: Right. But I think you need to be careful in your reference to SONET. John was talking about characteristics of the router network. If I look at transport network as a common facility network for all of our services, it's the lines that connect the boxes. What John is saying is that for data traffic those lines are bigger because the data traffic is already being aggregated is packets. The voice networks say connect up my 1.55 megabit T-1s. But the data networks say connect my 622 megabit OC-12 s and more frequently now my 2.5 gigabit OC48s or my 10 gigabit OC 192's.

COOK Report: But by the time the voice network gets on your intercity optical backbones, it has been packetized and becomes part of the data network packet streams, yes?

Afferton: It is sharing the common facility but it is not necessarily packetized and part of the same packet flows as those of the data network. Of course, Voice over IP may change that, but we must also consider our embedded facilities.

Role of SONET as Multiplexer in a Layered Network Structure

Strand: You need to think of the network as a layered network where the actual fiber is at the physical bottom. DWDM systems form the next layer and provide high capacity pipes that light the fiber. Sitting above the fiber and DWDM systems are equipment used to connect the pipes together to connect customers. This equipment includes optical cross-connects and add/drop multiplexers. This is the point at which there is a division between very large pipes for IP traffic and much smaller pipes with voice traffic coming from the legacy world. The smaller pipes might be DS3s which are mapped into SONET STS-1s.

COOK Report: Do these small pipes get mapped into some larger OC kind of a circuit?

Afferton: The answer to your question is

yes. This is the role of SONET multiplexing. There is some perception out there in SONET is bad and that SONET is incompatible with data. There are two aspects to SONET. First SONET is a mechanism by which you can provide framing for a signal so that you can transport that signal from one place to another, monitoring it as you do so. The second aspect if SONET is that the framing also provides a structure for aggregating traffic into larger bundles, i.e. multiplexing.

A router today can have a packet over SONET Interface and you can use Packet over SONET technology to put the packets inside of the SONET payload. For example, in AT&T's IP backbone between our core routers we typically have OC 48 and OC-192 facilities. They are SONET framed but they do not go through SONET multiplexers. They would be carried directly onto our optical backbone without having to be multiplexed with other traffic because they are already a large bundle. Voice traffic, and also private line and other services that don't require such large pipes, also need SONET framing but at lower rates. In order to carry these cost effectively through out our network, we combine them together with multiplexers. There is the SONET OC " x " hierarchy that facilitates this. When people talk about SONET being bad for data because it is inefficient or anything else, what they're saying is that " look we already have a very large pipe coming out of our router, you don't have any need for it to be multiplexed." So the objection is more to multiplexing than to framing.

COOK Report: And by lower rates you mean smaller bandwidth?

Afferton: Correct. People use the term IP over WDM or IP over glass implying that they do not need any multiplexing because the pipes are already large enough to be placed directly over the optics. And that is what we do in our network with packet data. We connect backbone routers directly to DWDM systems without SONET multiplexing. We provide this capability to external customers as well, which some people call "wavelength services".

We are also expanding this capability to provide customers the ability to interface to us using Ethernet, where even the SONET framing is dropped at the customer interface. However, depending on the size of the connection, we still may carry that traffic internal to our network inside of SONET. We have found, when you consider the customer's needs for protection and bundling with other services, Ethernet over SONET solutions are actually extremely cost-effective.

Strand: Even some of the packet carrying pipes can benefit from the SONET multiplexing by making them into big enough bundles to make it cost-effective to send them over the optical network. Today lots of access routers are connected by OC-3s and OC-12s and not OC-48s. Sixteen of those OC-3s could be multiplexed together into one OC-48. Think about this as is a stack with two converging flows. There is a TDM multiplexed flow coming from the legacy world. And then there's the tributary coming in there of OC-3s and OC-12s from the IPworld. Increasingly the traffic from the IPworld is merging in with this other flow from the legacy voice world right at the optical layer with OC 48s or OC-192s. Tom and I, as we sit here worrying about the plumbing in this whole thing, see all this stuff coming at us as high bit rate, normally SONET framed flows, of the OC-12, OC 48, and OC-192 variety.

COOK Report: Can what you've just described be seen as an effort to figure out how a carrier like AT&T can take its own voice traffic and the investment supporting that and map this with maximum cost effectiveness on to its new optical network?

Strand: If you go back to our image of the "v" shaped flow you will find that the optical network is largely agnostic to the question of IP versus voice. Investment in the optical network has benefited our ability to handle both of those flows. Small bit rate, private line and voice services are basically evolving at a slower rate. It is still growing and it is spinning off lots of revenue which we like. But the growth of bandwidth is driven by OC-3, OC-12, OC 48, and OC-192 and Ethernet point to point connections be-

tween routers. Consequently, these are the things that we look have to look at more and more as the drivers for how we design the optical network.

Afferton: Fortunately it's not like we have to build a different optical network to support both the new and the old. All the carriers are indeed seeing a shift in demand from voice to data. We are also seeing a shift in terms of demand which means that the bulk of data that is hitting our network comes there in larger and larger pipes.

COOK Report: When you say at the network in terms of larger aggregate levels do you mean right from the level of the enterprise customers on up or from the levels of Metro networks when they connect to the backbone? Or probably both?

Afferton: It is both. We have customers, ISPs, who come to us and would like us to build their backbone for them and would like us to provide them with a nationwide network of OC 48 pipes. We did this for @Home in 1999.

Fiber Backbone Deployment

We are also building the pipes for our own data network by means of which AT&T provides data transport services, VPN, or what ever else is needed to business customers. One of the decisions that we've made is that we will be facilities based. We are building out our infrastructure to handle our own services. We have been effectively doubling the capacity of our transport network year after year for the last several years.

We are in the midst of an aggressive fiber expansion program that we announced at the end of 1999. AT&T today has 53,000 route miles of fiber in our network. This includes of both of our long-haul and of our metro networks. Note of course that I am talking route miles. There are numerous fiber **strands** in each route mile. This gives us not only nationwide coverage, but a lot of penetration into a lot of cities. This is our imbedded network with which we've done very well.

However, looking ahead in anticipation of the growth of data services, we made a decision at the end of 1999 to expand our fiber by putting in another 16,500 route miles of new fiber. It is being installed predominantly along right of ways that we already have but it is also a nationwide rollout that picks up 30 of the top cities where traffic is originated and terminated. We are doing this construction with partners. We are sharing the construction costs and coming away from the project with fiber cables and spare conduits which will allow us to easily pull in new fiber at some point in the future. We have therefore made the investment to continue to expand the glass that we have in the ground. Also as John mentioned there is another investment in equipment that lights this fiber. These are our DWDM systems. We have made the investment to increase the capacity of these as well and to do this not just with our new fiber but also with our existing fiber.

Because we're making the necessary investments in our own infrastructure we are able to continue to build backbones for our own service networks; to continue to handle our legacy traffic; and to continue our own backbones of for IP and ATM and frame relay and to provide services for other companies as well.

COOK Report: How would you describe in a general sense the areas of functionality that are needed for network design in the industry as a whole?

Strand: Traffic patterns are very important. Voice traffic tends to be very short and the probability of two people chosen randomly wanting to talk at the same time decreases very rapidly as the distance between them increases. In the data world none of the patterns that determine voice traffic hold true. Initially in the voice world you had a hierarchy a lot like the tier one, tier 2, and tier 3 ISPs. Traffic was either local or toll within a regional area, or sectional, or national. As traffic built up there increasingly was enough traffic going from one regional area to another to justify building trunks directly between them without having to carry the traffic up from a toll switch in one city to a higher layer toll switch in a

distant city and back down to a another toll switch in the originating region. Over the course of time, as traffic built up, the hierarchy was flattened and the higher layer switches eventually went away. The equivalent of Internet tier 1 and tier twos went away and we ended up with a flat network.

Data Exchange is Less and Less Hierarchical

The analogy here is as the Internet grows and you start getting more and more traffic between lower-level ISPs, you have enough traffic to exchange directly between them without having to go up to an exchange point on a national backbone and back down again.

Something that is very important in looking at data traffic patterns is that if you have enough demand between two points, you can connect them directly with optical transport facilities. Every time you have to send you data through an intermediate router you face a situation where a line card on a router costs orders of magnitude more than a line card on an optical cross connect.

COOK Report: Are you saying then that, if I am in Philadelphia and I have access to a line card in an optical cross connect, and if I have a friend in Buffalo with access to another line card in another optical cross connect, and we wish to communicate with each other, that we could configure our line cards to talk directly to each other? Moreover that we could do this with out ever using the IP layer as a part of our communication?

Strand: Supposing your Philadelphia office is situated near an office that has an AT&T optical cross connect. You could then run a cable from a port on the router to the optical cross connect. To make this economical, of course, you need to want to send a lot of data to your friend.

COOK Report: And an optical cross connect is an intelligent optical device that would permit the provisioning of a data path at the optical level from one cross connect to the other?

Afferton: Yes. That connection could be at a variety of other bandwidth levels. It could be as low as a DS3 or as high as an OC-192. When you say cross connect, it implies that, at a fundamental level, what it does is make a connection from one port in a box to another port. That is the physical aspect. I can take a signal coming from this port and connect it to another port. What you want is to have a network of cross connects that can communicate with each other and act in an intelligent way to establish a connection from a port in a device at one location on your network across the network itself to a port in a device at a distant location.

You also need to be careful about your use of the term optical. The term has been loosely applied and means different things to different people. The term optical cross connect has generally been used to describe a cross connect that can make connections between two optical ports in an intelligent device that has some built in switching or transport capability. And that is also to say ports with optical signals coming in. Now there are some products that make these connections by means of electronics. Your optical signal is put through electronics and then converted back into an optical signal. This kind of conversion may be important because it can allow you to rearrange the signal at a higher level of granularity. To give one example of what I'm talking about, you could rearrange 45 megabit signals inside of a larger pipe.

Introduction of Software Control at Optical Cross Connects

But there are other products out there that take the light coming in and direct it from one port to another. They do this entirely with optics with no examination of the constituent signals. There are vendors out there who will debate with each other which cross connect solution is better and more cost-effective. Our concern is much more with the functionality offered than the technology by which the functions are achieved. These examples of customer provisioning and optical peering are all things that we see coming down the road with varied service availability.

The underlying technologies to enable these capabilities for our customers are what we're spending a lot of time working on and getting ready to roll out.

Strand: Just picking up where Tom left off then, in looking at your Philadelphia to Buffalo example, we have these two routers each connected to optical cross connects. Imagine a mesh network or a grid of DWDM multi wavelength optical transport systems. And at each node in this grid you have one of these optical cross connects positioned. As a result, if you want to connect from one router to another, if they are on the network with the optical cross connects, you have to tell each optical cross connect: connect port "i" coming in to port "j" going out. In turn, port "j" going out is hard wired to a specific wavelength on one of these DWDM systems. It goes on to the next hop and then you do the whole thing over again. This is basically the way the optical network transmits data from one node to another. When people talk about intelligent optical networks what they mean is the ability to reconfigure by means of software this node to node transfer process very rapidly. Along with increasing capacity and reducing unit costs, the leading edge of development in the optical network at this point is the introduction of more and more software control.

In this general context then we must look at our customer needs and, in keeping with available technology, decide what type of architecture will best and most cost effectively satisfy them.

Afferton: While in a voice network everyone can connect to us with a 64 kilobits per second DS0, in a data network we have to have the flexibility to support a whole variety of interfaces that customers are going to want to use. If they want to interconnect in order to complete a LAN or MAN or WAN and want to use only a native Ethernet interface, we want to have the ability to do that for them. If they want to interconnect with an OC-192, we need to have the ability to offer them that.

Strand: Because data is not inherently symmetrical the way voice is, this opens up a whole other set of things than one

might conceivably do in an optical network.

More Rapid Provisioning

More rapid provisioning is one that has been driving the Optical Internetworking Forum and other organizations like ODSI, ITU and the IETF. While use of the voice network changes very rapidly on a moment to moment basis, there is a very well understood architecture that is managed very efficiently by the use of traditional voice circuit switches. What is different about these large pipes is that they are down at the transport layer in a part of the network where we are not all used to putting up and tearing down gigabit size connections the way that we do with voice connections.

COOK Report: So what you're looking for is an architecture for the data traffic that uses your resources with maximum effectiveness and meets the full range of your customer's expectations?

Strand: We are trying to establish an architecture that can do it in many ways better than we were previously able to do it.

Afferton: I think what we are seeing is a shift to the point where the data is more the rule in the network than the exception. The presence of data preponderant traffic causes you to want to have a more dynamic transport network. To have a level of churn similar to that in a voice network but at a much larger scale of bandwidth.

Strand: In this context the technology is starting to give us the capability of reconfiguring these big pipes very rapidly. We are examining how to most effectively apply this capability in our own network. Looking at current situation, we see that the average utilization rate of a data private line is very low. According to Odlyzko's current study it is approximately 10%. This means that 90 percent of the bits are unused. A large chunk of this is due to time of day and day of week considerations because businesses, of course, close down at night.

In the future there should be the potential of not making a customer pay for the pipe during that time that he is closed in the evening and giving it back to him again when his business day begins. Another example of why rapid provisioning is needed was pointed out by Avi Friedman in his interview with you two months ago when he complained that his customers had to wait to six months to get some voice circuits provisioned. If we can do rapid provisioning, both across our metro networks and core, we can solve that problem.

Afferton: While rapid provisioning is certainly a need, extending that rapid provisioning from the metro across the wide footprint is critical to having it be a meaningful service for a customer. We believe that it is very important to have capabilities to extend auto provisioning throughout metro networks, intercity networks, and access networks. We are well positioned both in our understanding of that problem and in AT&T having the comprehensive assets to take advantage of the solutions.

COOK Report: I see that the Optical Internetworking Forum came up with a list of ideas of services that would be dependent on rapid provisioning.

Afferton: Indeed the OIF carrier group is a good place for people to brainstorm ideas.

COOK Report: These are service ideas and not commandments chiseled in stone?

Afferton: Correct. One idea was that rapid provisioning could be done internal to a carrier's network. Under carrier control we could put up circuits faster in response to customer requests.

Working on rapid provisioning just within our own network will be something of a first step and it will be one we take because we will not have to worry about issues of security and billing with external customers yet. The second goal of bandwidth-on-demand might occur when a customer wired a fixed port to the edge of the optical cloud and then not have

that port connected through the network unless they wanted to be. In this application the customer may be subscribed to the optical network but only pay for usage when it occurs. For example, perhaps I was a customer of storage area networks and wanted to do backups for a period of time every night. Under the third idea, an optical VPN, a pool of bandwidth would be given to a customer for use among of a bunch of specified ports. The customer is given a subset of the network's capacity which no one else can use. Within a subset of capacity the customer can arrange its network however it wants.

All of these right now are general descriptions of potential services. However we're working on a network architecture that should permit them to become reality.

COOK Report: And each carrier at some point in the future when these ideas are implemented for the first time will likely do so in its own unique way? You would have your own unique flavor while another service provider might do it a bit differently?

Strand: Yes. We certainly will because we all have an interest in differentiating our own products.

Now, if you look at the overall costs of running the network, you will find that the intercity long distance cost is small compared to the cost of access and the last mile. One of the things that is very attractive to the carriers in the OIF carrier group is that this intelligent optical network structure, these software controlled optical cross connects potentially can reduce a lot of the costs of actually running the network. From an architectural point of view the second message here is that you make a big mistake if you just look at capital costs when you are trying to decide what type of architecture is attractive.

Afferton: In making decisions about our network architecture, the choice of technology that will allow us to simplify our operations and lower their costs is also important.

The AT&T Optical Transport Network

COOK Report: So how is AT&T solving these problems?

Afferton: There are a couple of different layers to our approach. I will go through them one by one and build up to the network view. I'll start with the bottom layer which is the glass in the ground. We base our offerings on our very extensive national fiber network which, as I have pointed out, we have also just upgraded.

We must, of course, light our fiber. To do this we have chosen DWDM systems both for our metro and intercity networks. Let me start by talking about the intercity segment. In talking about this segment of our network I try to use intercity rather than long-haul as more of a generic term since long haul is often applied mainly to voice. We are currently deploying WDM systems that are capable of handling up to 400 gigabits of capacity on a single fiber.

The systems that we are deploying today we are dual sourcing from Lucent and NEC. These systems are capable of supporting 2.5 Gigabit OC-48 signals or OC-192 ten gigabit signals. In addition, looking at the continued rise in bandwidth demand, we made a decision last year to pursue a system that is capable of carrying over a terabit of capacity on a single fiber. As a result of a public request for proposal process, we selected the Spectral Wave 160 platform of NEC. This will support 1.6 terabits bit of capacity .

COOK Report: Would you put that in terms of lambdas?

Afferton: The older systems will support forty of the ten gigabit lambdas and the newer 160. Certainly NEC's not the only vendor out there and we are not the only company pursuing terabit capable systems. But, in terms of our network needs, we are very comfortable with the choice we have made.

COOK Report: How will you be using the higher capacity boxes?

Afferton: We have the existing systems installed and supporting existing traffic. As we continue to grow, we will certainly fill up the wavelengths on those systems as appropriate. We will also light the new systems, for example, on the new fiber, as well as on the old fiber in places where we need capacity. In fact with the NEC equipment, we can take out part of the system that belonged to the previous generation and upgrade the remainder of it with the newer technology. The 1.6 terabit systems are in our lab now and we will be rolling them out live into the network later this year. We feel very good about these both in terms of unit costs and in terms of the intercity capacity they will enable us to provision.

COOK Report: Are you looking at this equipment to sell lambdas?

Afferton: When you talk about selling lambdas, it really means you're giving the customer the capacity equal to one of wavelengths that you carry across the network.

COOK Report: Right. From point A to Point B.

Selling Lambdas

Afferton: We were among the first to do this and have been doing it on a wide scale since the beginning of 1999. I mentioned @Home earlier. There was one large-scale example where @Home came to us in the early 1999 and said: "we want to have someone build us a nationwide set of OC 48 pipes." We told them that we could use our existing infrastructure to provide that.

There is a terminology problem floating around. We use the example that we're providing an OC 48 private line service or an OC-192 private line service. In addition to the @Home example, where we were one of the first to provide a nationwide network for an ISP, we were the first to provide an OC-192 private line service, in other words a "lambda or wavelength service" at 10 gigabits at the end of 1999. Now we offer OC-192 private line service on more of a national scale.

In order to offer such a service you need to be able to connect your customer to the WDM system and do it in such a way that you can clearly monitor their traffic in order to show that you are meeting your service level agreement and in order to be able to isolate troubles. To do this is where the transponders come into play. Transponders act a kind of gateway into an optical transport system. They take a "plain vanilla", standard optical signal and convert it through electronics into an optical signal that has the right wavelength and other properties to carry it through the DWDM systems. It is important to have your customer connect to a transponder because, it gives you a very clean controlled signal that you can work specs off of to hand to the customer. It provides a demarcation point into your network.

Because AT&T had already deployed transponder based systems, we were able to deliver those services based on our architecture which inherently supported them. There are other carriers out there who were using proprietary systems where you did not have transponders . Whatever equipment that was connected to the WDM systems had the specific wave lengths built into that client equipment. When you did something like this, you would need another special box to enable those customers to connect directly to your WDM system.

COOK Report: And that special box would be what?

Afferton: It could be some kind of an adapter. There were companies that started making those. There were others carriers where, in order to get onto a WDM system, you had to go through a SONET multiplexer. That was not cost-effective compared to directly connecting to a DWDM system. Providing private lines directly off a DWDM system, which people would call a lambda service, is something that is inherent to a transponder-based network.

COOK Report: It seems that what you're talking about is an example of why standards are needed?

Afferton: This is at the transmission

level. There, the standard required is well established - normally just a standard short reach laser. And the initial version of these lambda services are still somewhat static. When you are connecting the customer to the WDM system, you are basically hard wiring them to it. The way to make these services more dynamic is through the introduction of intelligent optical switches which is the next and third layer of our architecture which I will get to next.

Now one other aspect of your DWDM system is its length. Some people said that they need ultra long-haul systems where they can reach long distances without having to regenerate their optical signal. At a very simple level people said: the further you go, the cheaper it is because you have fewer intermediate regeneration points. But I think this is a gross oversimplification because obviously, if you go further, you need to improve your optics and add to the cost of the end point in order to achieve those longer distances. We have been asked before how we feel about ultra long-haul systems and my response is that I'm not particularly obsessed about being able to say we have the longest distance. What I am obsessed with is being ultra cheap by means of whatever technology will enable us to do that. In the end when we want the systems that will give us the lowest cost per bit per mile, while also taking into account operational costs.

Metro Network - Connecting the Plumbing

Let me turn now to what we're doing in the metro network where we're not aggregating as much data together. There your bandwidth demands are going to be somewhat less. But, as we pointed out earlier, there are customers out there who want to have a full wavelength carrying a large amount of bandwidth handed off to them. Therefore, you still need the capability to manage and deliver large chunks of bandwidth in a Metro Network. For this purpose AT&T is deploying, on a more selective basis where customer needs can justify it, Metro DWDM systems. Those systems are capable right

now of deploying 24 to 32 channels of 2.5 gigabits on a fiber. We will be rolling out later this year support for 10 gigabits as well.

In the metro network, you run into a situation where people want other interfaces besides SONET OC-48s or 192s. They might want full gigabit Ethernet for example between two places. So this can be another flavor of what people would call a wavelength service because they want a whole pipe delivered to one place or another and that pipe and fills up an entire wavelength on a DWDM system. Consequently the systems we are deploying have the capability to support protocol's beyond SONET. They can support Ethernet, for example, and they have the flexibility to support other protocols as well .

Going back to the definition of wavelength services, if you're talking about the flexibility to support something besides OC 48 or OC-192 SONET, the answer is yes. That is another aspect of our architecture which is particularly relevant in the metro networks.

That's what I will call the optical layer. We have the glass and the ground and now we have lit it in the metro and intercity network with DWDM systems. There are some services that might directly interface with those networks. But for the most part what we have is the plumbing. We have a lot of that plumbing. As a result the question becomes how do we quickly and intelligently connect that together for end-to-end services to customers. That is where the introduction of intelligent optical switches becomes relevant.

COOK Report: And your "plumbing" becomes a question of how you connect your metro and long-haul circuits?

Afferton: Correct and this is where intelligent optical switches come into play. We take all these WDM systems and connect them together through intelligent optical switches. This goes back to the comments I made earlier about the optical cross connect. Again it is not so much the technology as the functionality with which we're concerned. The functionali-

ty we are looking for is the ability to be able to connect together all these WDM systems with their very large capacity and then have full visibility into this inventory of that capacity. As we wire up the WDM network to the switches, we want all of the switches talk to each other and figure out how everything is wired up. And then our operations staff can say "a hah, I have this pool of bandwidth available to me." We do not have to make these calculations by accessing some offline database. Rather the term that we use is that the network is the database of record.

Because of the switches' ability to communicate with each other, the network also has the intelligence to establish a circuit by connecting all these links together. I merely have to specify a start point and an end point. At that point the network nodes will communicate with each other and put up the circuit I've requested. It will do it, end-to-end, using what really are IP routing techniques. Once you do this, you have the capability to rapidly provision circuits. You can quickly put up and tear down circuits and, as a result, you have a more dynamic network.

COOK Report: What precisely did you have in place at this point in time?

Afferton: Right now we're in the final stages of testing these boxes. We will start soon with a small footprint to make sure we test it well in the field environment. After that it will grow out to a nationwide footprint. It is targeted to have nationwide coverage toward the end of the year. It is a fundamental point of our architecture to which we are fully committed. As far as the exact timing of the rollout we know that the vendor performance plays into that as well. We have not said publicly who the switch vendor is.

COOK Report: Are you the first in the industry to do this? In is anyone else doing it even in a small footprint?

Afferton: I don't want to try to make a guess on that subject because there are others out there who are talking about doing similar things.

COOK Report: But you are not aware of another service who is doing this even in a small footprint at this point?

Afferton: It is hard to interpret what other people are saying about their networks.

We are committed to our own efforts. We expect our customers to quickly see the benefits of rolling this network out. Initially, the capabilities of rapid provisioning will be used internally. Our own operations folks will quickly turn up new circuits in response to bandwidth needs of our internal networks and our external customers. But on day one, we're not going to allow customer to directly control these switches. This network will provide the platform to enable bandwidth on demand and virtual private network services down the road, when appropriate security, billing and other complementary capabilities are in place.

COOK Report: Is it safe to assume that this provides the ability for you to handle your own infrastructure more cost-effectively?

Afferton: Correct, and this is the initial benefit. In order to maximize these benefits, we must also consider metro networks as well. We are extending this kind of capability out into the metro networks through the use of what some folks call multi-service provisioning platforms. These are basically SONET multiplexers that have added intelligence and can handle not only SONET TDM type traffic but also, for example, can map Ethernet into SONET, and can do some ATM switching and so forth.

We have started rolling these out into our metro networks in order to reduce our costs, as far as SONET traffic is concerned, but also to have the intelligence to inter-operate with the intelligent optical switches so that we can then do this automated provisioning end-to-end. For example, the purpose is to go from an enterprise and then a metro network across an intercity network into a metro network on the receiving end and to wind up at another enterprise.

Inter-operability

Inter-operability will come over time. This is really the focus of the standards organizations. For example, we want to ultimately be able to have intelligent protocols that can do automated provisioning and even coordinated restoration between the clouds of equipment from different vendors.

COOK Report: And the reason for connecting to different clouds is that you might have to connect AT&T customers with customers from other service providers?

Afferton: More to the point within our own network we would have vendor A building one metro network, and then vendor B building our intercity network, and then vendor C 's equipment in a second Metro Network. This is the first and most immediate application for the standard interface. Therefore a big part of our standards effort is to develop these protocols to enable the inter-working between the equipment of different vendors. This is really the essence of the architecture that we're rolling out right now.

COOK Report: Are the benefits of these protocols likely to be available say three months from now or will be more like 30 months?

Afferton: It is hard for me to say. There are different aspects to the intelligent optical routing protocols. First there is the interior gateway protocol running within a cloud. Because these protocols are running among a set of boxes that come from one vendor they initially can function on a proprietary basis. There is an effort underway in the IETF to create a standardized framework for these interior protocols, called Generalized Multi-Protocol Label Switching (GMPLS).

Then there is a connection between two clouds — a network to network interface. Work is starting on this area in the IETF, in the ITU and also the OIF. There is also another interface between the optical network and the underlying data network. This is where instead of a person calling and saying give me this pipe for two hours the equipment itself requests the

bandwidth. So when routers see congestion, a router can say to the optical network I need another connection fired up between myself and this other router. That is a user to network interface (UNI). The OIF has been largely focused on this interface.

In fact at Supercomm next month (June 2001) what I will call the alpha version of the UNI being developed in the OIF will be demonstrated among a large group of different data equipment vendors and optical equipment vendors. This is a promising first step towards multi-vendor interoperability among dynamic optical networks.

Strand: Before you can do any trading in bandwidth as a commodity you really have to have these capabilities in place? Yes?

Afferton: Yes. Within a network like AT&T's you have to talk amongst the various parts before it makes sense to talk about the more general kind of stuff between different carriers.

From Corning a Much Cheaper Method of Laying Fiber

On June 7, David Isenberg <<http://isen.com/>> wrote in Smart letter #56:

The act of digging up the streets is becoming passé, if not obsolete. Corning MCS-Road cable packs up to 144 fibers in a 7mm diameter cable. You "install" it by cutting an 8cm slit in the pavement with a circular saw. You kick in the cable, tamp a rubber strip on top and seal the slit with goo. It is up to eight times faster and five times cheaper per route mile than trenching according to Corning. . . . The Corning MCS (Micro-Cabling System) product line expands beyond MCS-Road. It includes two different ways to pull cables through drains and sewers, plus cheap and easy splicing, interconnect and repair systems. Get more details at <http://www.corning-cable-systems.de/en/products/mcs>

Level 3 Wants to Be Global Carrier's Carrier

An Introduction to its Structure and Lines of Business

Editor's Introduction and Note:

Readers should treat the following interview with as an Introduction to Level 3 which is sitting on an impressive infrastructure of global proportions and has cast a very interesting business model in view of what some see as the doubtful ability of green field players to handle debt in what looks like a glut of fiber. Although the talk with Morely gave a lot of information, we felt at the end a comprehensive picture was lacking. Consequently, we dove back in and have now completed two more interviews. One was with Ron Vidal is Group Vice President, New Ventures and Investor Relations and the second with Robert Hagens, Senior Vice President Global Architecture. Both are very extensive and together they flesh out the framework that follows in fascinating detail. The draft of Vidal is 8800 words and we estimate that Hagens will be about 7000. Unfortunately readers will have to wait until the next issue (September - to be published early July) to read them.

Meanwhile, Andrew Morley is Sr. Vice President of Global Strategy for Level 3 Communications, Inc. Andrew joined Level 3 in October 1997. Prior to his current position, Andrew was Sr. Vice President of Global Transport and IP Services. He has also held senior positions in Marketing and Corporate Development. Prior to joining Level 3, Andrew was a Partner with Marquette Venture Partners, a venture capital investment firm, where he invested in emerging information technology and communications companies. Andrew has an MBA from Harvard Business School and an Engineering Degree from Yale University. We interviewed Andrew in two sessions. One on May 9 and one on May 10.

COOK Report: Would you start with a quick overview of your U.S. infrastructure and then talk about how you are provisioning it?

Morley: In the U.S. we have a 16,000 route-mile twelve conduit intercity network serving 54 markets. We have 26 U.S. cities with local fiber networks. We have approximately 6 million sq. ft. of data center and technical space globally. We have pulled fiber in the first conduit and will have migrated our network traffic that was traveling on leased lines to our own fiber by July 1 of this year.

COOK Report: What has happened to the XO-owned portion of your fiber?

Morley: XO had purchased 24 strands on our U.S. network. We have recently announced a deal with XO where they will keep the fiber, but they are not going to light that network. Instead they have announced that they will be purchasing wavelength services from Level 3, which will serve as their underlying transport network and we will be buying some of the optronics that they were in the process of deploying.

Building a Global Infrastructure

COOK Report: Who has been buying your dark fiber?

Morley: We have dark fiber customers, ranging from Cable and Wireless to France Telecom who have purchased over \$4 billion in dark fiber and related services.

COOK Report: Where you go in the rest of the globe?

Morley: We have two intercity fiber rings operational in Europe totaling 5,300 km in length and serving nine major markets. Each of the rings has between six and nine upgradable conduits. We built 12 originally and have swapped some of them with COLT, with whom we did the joint build. In addition, we have what we call our yellow system - a 1.2 Terabit trans Atlantic cable that we built and on which we have sold a substantial amount of capacity to Viatel and

Global Crossing. We also have capacity on AC-1 as well as TAT 14 cables. We own 11% of the capacity of the Japan U.S. Pacific cable. This is a consortium-based cable. We recently announced a partnership with FLAG for our Pan Asian cable. We are completing a cable between Hong Kong and Tokyo and building a subsequent cable that goes to Taiwan and up to Korea and then back to Japan. Investment in this cable is being split 50/50 between FLAG and Level 3.

COOK Report: Other than international links what kind of business are you doing in Asia?

Morley: In both Hong Kong and Tokyo we have substantial Gateway facilities. We have back haul facilities in each of these markets and interconnections with all significant players in them. This allows us to offer a full range of transport, IP, colocation and softswitch services in each of those markets.

Now we currently have co-location and distribution facilities in 54 markets within the United States. In nearly half of those markets (26) we have also built our own upgradeable metro networks. If you combine our long-haul network with our metro networks and co-location space you will have what, from our point of view, are the key elements in offering the lowest possible end-to-end cost to our customers.

Voice in the Data Mix

COOK Report: You are certainly building for data networking, but what about voice?

Morley: We consider ourselves a carriers carrier and are focused on helping carriers outsource their underlying networks. We also work with ILECs, CLECs, PTTs, Wireless providers, very large ISPs and cable providers when they're looking at outsourcing. Of course we also run a very large IP network and therefore work with content providers.

In the voice area we have a Softswitch platform, which is a voice-over IP platform. We are today doing 8 billion minutes a month of traffic over that platform with our managed modem and voice services. [Editor's note: according to <http://www.level3.com/us/news/newsreleases/1,1345,2001Mar21-5379,00.html> "Pioneered and co-developed by Level 3, a softswitch is computer software that emulates the functions of traditional voice circuit switches to control and process calls over a communications network. Similar to a computer operating system, softswitches feature an open platform on which developers and Level 3 customers can build innovative and specialized applications that make communications less expensive and more productive."]

Softswitch supports managed modem service for the AOLs and Earthlinks of the world. Softswitch is our platform for providing voice and other switched IP Services. And Softswitch does provide for the interface between our IP network and the public switched telephone network. However Softswitch manages flows across our IP network as well. We also use Softswitch both to direct voice-over IP and switch IP data across of your network. We have been using Softswitch for our data services since late 1998 and for our voice services since late in 99.

COOK Report: Are you providing managed SONET services for people or must they provide it for themselves if they need it? Also what about gigabit Ethernet?

Morley: We sell a whole range of services: everything from dark fiber to wavelengths to SONET to IP to voice. For those customers that want to buy them, we do sell traditional, protected SONET based optical services.

COOK Report: Are the SONET based services available in the U.S. over virtually your entire network?

Morley: Yes. We have been selling them since really early in 1998 over a leased line network. They are available over our constructed network today in the US and Europe as well as over our Atlantic and

Pacific cables. In addition to those optical services, we sell IP services with 100 megabit and one gigabit per second ethernet interfaces, as well as SONET (OC-x/SDH) interfaces.

COOK Report: If I want an OC pipe that runs IP, I can get it from you?

Morley: Yes. Everything from DS 3, to OC 3, to OC-12, to OC 48 IP Services. Last year we introduced what we call EPOP. This basically provides metro Ethernet services to access our IP network. It provides a very very low cost alternative to the ILEC local loop to extend Ethernet interfaces out to our customer premises.

IP Backbone

COOK Report: Are you selling transit across your IP backbone? If I am a local ISP and am selling connectivity to the global internet to my customers, will your backbone be all that I need for delivery of my customer's traffic?

In other words, do you guarantee to get my traffic were needs to go as long as I pay your connectivity charges?

Morley: Absolutely. We are tier one backbone provider and have Substantial peering and interconnection with all the other tier one Providers.

COOK Report: Are you transit free? In other words do you pay other networks any fees to deliver your traffic? UUNET is very often the backbone to which it is nearly impossible to get transit free peering. Are you fully peered and transit free with UUNET?

Morley: Yes. We have settlement-free peering with every major IP network provider in the U.S. The vast majority of which is done through private peering points, rather than public peering points. This ensures the optimal network performance through these interconnects.

COOK Report: Does Level 3 have any enunciated positions on issues of settlements and paid or transit free peering in its IP networks?

Morley: Given the traffic volumes, scal-

ing in the interconnections is a key issue. We are generally looking at multiple OC-12 or higher interconnects with our peers in multiple locations. The key in being able to scale is to build substantial physical infrastructure to our peering partners in the appropriate location points to allow us to scale capacity quickly and efficiently.

SONET versus Ethernet Network Architecture

COOK Report: Is gigabit Ethernet being used in interconnections anywhere? Or is it all SONET?

Morley: It is predominantly SONET. There is some discussion of moving toward using Ethernet which may make sense for such interconnects. Historically, the interconnects have been SONET.

COOK Report: And as far as moving to Ethernet is concerned, is all that is required for you to get the appropriate people at a co-lo point and install switching equipment?

Morley: That is one possibility. Another possibility is that our metro fiber networks would allow us to do that with our peering partners without them being all co-located.

COOK Report: To do this at co-location facilities, would you be using a technology similar to that described by LayerOne where the bandwidth provider takes a Ciena Core Director and uses it to groom output circuits for its customers? Is this a part of your overall picture?

Morley: Having that kind of capability and functionality is a key part of the picture of where Level 3 is going in order to serve its customers. There are a number of efforts under way to move toward more scalable interconnection architectures and processes.

COOK Report: If you look at how you are going to join your wide area and metro area networks, presumably the Ethernet EPOP is the critical component for achieving this?

Morley: The MAN WAN interconnection spans several layers. First would be the fiber layer. Of course we also consider our metro networks to be a key low-cost enablers in getting customers onto our backbone networks. In the fiber layer we have very high fiber count, multiple conduit, metro networks. We will allow a customer to take dark fiber from our colo facility out through our metro market, across our backbone and into another metro market. At the wavelength layer we also use our metro fiber to extend the metro wavelengths again from our facility to the particular location where the customer would like to pick up that wavelength. From a private line standpoint, we again use our metro infrastructure. We put our SONET optronics on our metro infrastructure to distribute traffic off of our long haul, again at the SONET layer, across our metro network to customer locations.

We can, in other words, deploy a SONET-based Metro infrastructure to deliver those SONET services from our long-haul network to our metro customers. Finally at the IP layer, either we can use the SONET infrastructure and run IP over it yielding SONET-based access to our IP networks of OC-3s and OC-12 s. Or we can offer direct Ethernet services. In this case, either gigabit Ethernet or 100 BaseT running directly over our fiber infrastructure and bypassing the SONET infrastructure at the metro layer.

COOK Report: How would you distinguish your metro Ethernet services compared to those of Yipes! and Telseon? Are they just plain vanilla with no frills?

Morley: First of all Yipes and Telseon are both customers of Level 3. Yipes! and Telseon are focused at the enterprise customer. Our focus at Level 3 is on supplying carriers (ILECs, CLECs, IXCs, DLECs and to the Yipes! and Telseons of the world. - namely companies who sell to enterprises).

COOK Report: And probably also Sigma?

Morley: Exactly. Sigma is also customer of ours. We are really focused on provid-

ing low-cost, high quality, high bandwidth services to what we call major aggregation points in a metro market. Our customers to whom we sell in these markets would take what they get from us and use it to provide services to enterprises within the metro markets.

Economic Analysis

COOK Report: let's move on to the gloomy economic analysis of Ravi Suria! He was fairly nice in what he said about Level 3. Nevertheless he thought you were still in trouble. What is going to carry you through?

Morley: Suria is making the assertion that the telecommunications market is overcapitalized with debt. A lot of money has flowed into the telecommunications industry. And if you look at things painted with broad brush strokes, that is true. As a result of this some enterprises have been funded that probably should not have been. But I think that is also very generalized view of the market and that you need to differentiate between those businesses with business models that give them unique advantages and those businesses that do not have such models. You need to dig a little to understand who the winners and losers are going to be.

Clearly it is a difficult time in the Telecom Industry. But Level 3 as a company has approximately four and one-half billion dollars of liquidity . We view this cash as a major competitive asset. We are fully funded to cash flow positive. In other words we have enough cash and liquidity to enable us to execute on our business plan as I've been laying out to you until our revenues are sufficient to cover all costs and capital investment at which point Level 3 will be a fully self funding entity.

COOK Report: Well someone must be talking about the commitment of time lines to accomplish this?

Morley: Yes. Current expectations are that we will reach EBIDTA positive on a run-rate basis by the end of this year and cash flow positive by late 2003.

COOK Report: You certainly have an impressive global infrastructure, but you also have a lot of debt and you are going to need some substantial income to pay the interest are you not? For example let's consider lambda sales. How important this is? How does it fit in what you are doing?

Morley: Those sales are a key part of our product portfolio.

COOK Report: With what equipment are you lighting your fiber network?

Technology Used

Morley: With equipment from Nortel.

COOK Report: What about routers for your IP network?

Morley: We're using Cisco and Juniper. We are using Juniper for MPLS and for the large OC-192 segments of our network. We use a mixture of Cisco on in the other segments of the network.

COOK Report: If you look at what you were selling at the end of the first half of 2001 how would you categorize your market? We've agreed that lambda's are an important part but before we get into them in detail what else is there?

Morley: We sell a range of services: fiber, SONET-based private line services, lambdas, as well as IP Services, soft switch services which are a combination of the voice over our IP infrastructure and managed modem services. These managed modem services are sold on a wholesale basis as banks of dial up modem platforms controlled by our soft switch infrastructure. We use our Softswitch platform to eliminate circuit switching in all of our network. The Softswitch platform controls the modem banks and then interfaces with the public switched telephone network.

In the traditional architecture, a carrier or an ISP puts a circuit switch in its network that talks into the ILEC, receives the call, and then passes the call off to a modem bank. By having our modem banks controlled by the Softswitch and interfacing

directly with the Public Switched Network, we can remove that circuit switch(5ESS) from our network. The 5ESS costs about one hundred dollars per port and with tens of thousands of ports per switch they can run into costs of millions of dollars per switch. We eliminate the need for this very expensive piece of equipment. Softswitch enables us to take advantage of price performance improvements that are happening in modems and IP equipment. Because it runs on a Unix or Linux platform it can take advantage of open source software. It can also serve as an open platform for others to build applications on top of.

COOK Report: Who are some of your customers and what would be some examples of the applications they are building?

Morley: AOL wholesale dial up access and Earthlink are two of our managed modem customers. They use us for the underlying infrastructure. It's our phone numbers, our hardware, and network. They are focused on providing service over our infrastructure.

COOK Report: Then local soft switch pops could be used to provide dial in services to multiple national providers? In other words the same pop would authenticate users of AOL, Earthlink and other ISPs using your service?

Morley: The modem ports can indeed be pooled to efficiently utilize the platform and having been authenticated people who dial in can be passed off to their respective service providers.

COOK Report: How about focusing in on wavelengths?

Morley: We do view wavelengths as a core service offering. Although it is a relatively new service for the industry, it is fast becoming a preferred service for carriers and others who are looking to purchase wavelengths, basically as a way of outsourcing their underlying network infrastructure.

Carriers who had considered themselves as owners of network assets are now beginning to talk to us about the purchase

of wavelengths as a way of avoiding much of the upfront capital equipment that they would otherwise have to make. (Be it trenching fiber, buying optronics or paying the salaries of the people they would need to run their own network.) They can augment their network in what could be looked at as a just-in-time basis.

COOK Report: So a lot of the potential customers for what you're talking in the would be a Sprint, an MCI, or AT&T?

Morley: Yes. Those kinds of carriers, large ISPs, regional carriers, and foreign carriers as well. For example France Telecom is both a private line and dark fiber customer of ours. They are not a wavelength customer yet but they would be logical kind of candidate for becoming one.

COOK Report: For example, if I am Sprint and my circuits, say from Atlanta to Houston and then Los Angeles are becoming crowded, I could come to you and inquire about buying a wavelength that would go from city to city to city as long as the cities I wanted to reach were on your network?

Morley: Indeed. As they need capacity they can buy it in much smaller building blocks than if they were to construct it themselves.

Lambdas and Unit Cost

COOK Report: Take me through the procurement process. If I am a carrier what happens when I call you? Take a wavelength for example. Can I buy one in various bandwidths?

Morley: We sell lambda's in bandwidths of either 2.5 gig or 10 gigabits.

COOK Report: How are they priced?

Morley: On a segment basis. From point A to point B on our network. It is distance related but it also varies depending on the segment and where you want to terminate the wavelength. There is a segment cost of that gets you, for example, from Chicago to New York. And there is also a termination cost depending on where you

would want to break out and distribute the wavelength.

COOK Report: When you speak about termination in New York City are you talking about terminating in either 60 Hudson Street or 111 Eighth Avenue?

Morley: To answer your question about the way it works: we can terminate that wavelength in our co-location facilities. The customer can pick it up at the facility and distribute elsewhere within the city. Or we will terminate the wavelength directly to the customer's facility. For example at our co-location facility the customer can put ADM boxes on that wavelength, and then distribute SONET services from the facility. We are in both 60 Hudson Street and 111 Eighth Avenue. We also have a 550,000 square foot facility right around the block from 60 Hudson Street. We use this facility to hand off traffic to our customers. We also serve our IP and Softswitch customers out of this facility. Otherwise we will use our metro fiber networks to deliver wavelengths directly to our customers' locations.

COOK Report: Would you be able to offer your customers like France Telecom attractive space at your own facility with a cost presumably enough below that of the 60 Hudson or 111 Eighth Avenue? Doing this might make it worthwhile for such a customer to connect at your facility knowing that it likely would need to be located in one of the other more public facilities as well.

Morley: Yes. The bottom line is we will distribute the wavelength to wherever the customer wants to take it. The customer may want to terminate in the co-lo facility and break the wavelength down into individual services that it is responsible for distributing. Or the customer may want to take that wavelength to their primary location. That location could be a facility like 60 Hudson or it could be a large pop at some other location.

COOK Report: Presumably it is safe to assume that you have your own metro fiber in both in New York and Chicago?

Morley: Yes. And this is a key point.

That metro fiber is a resource in providing Network outsourcing for these networks. It is one that allows us to very cost-effectively distribute that lambda from our facility all the way to the customer's premises.

If you compare us to our competitors, we believe that you will find that we have a much lower cost structure, a much more robust service offering, and a new technique that will allow us to guarantee to provision lambdas in less than 30 days. This is a provisioning interval that is a fraction of what is now standard in the industry.

Our metro networks of course allow us to distribute wavelengths to the customer location. By having large fiber counts in the metro areas, we can distribute the lambda's without having to utilize DWDM which would add significant cost to terminating those services whether they were lambda services or SONET services. We believe that this combination of long haul and metro facilities gives us the lowest cost structure in the industry today.

Now our network is also very importantly an upgradable network. What we mean by this is that our network will allow us to continue to deploy new technologies very rapidly and to take advantage of price performance improvements that are occurring not only at the optical layer but also at the IP layer.

COOK Report: But what would be an example of a fiber network that would not be easily up-gradable?

Morley: By way of example we built 12 conduits across our U.S. network, and six to nine conduits across our European network. Today only one conduit is filled. The unit cost for the transmission of data in an optical network is heavily driven by the electronics and the capacity of the electronics that you deploy on your fiber. The ability to use new generations of fiber over time will allow you to take advantage of new generations of optronics. New optronics can be optimized with new fiber to give yourself better unit cost than if you find yourself forced to deploy new optronics on old fiber.

COOK Report: So you are talking your ability to pull new fiber through a second conduit and to light that fiber instead of the first conduit first-generation fiber?

Morley: Yes. And we have just completed pulling the newest generation of 10 gigabit fiber through conduit number one in our network. We have deployed the newest technology and that has given us a cost advantage over such networks like Broadwing and Qwest that recently deployed a previous generation. And we expect that a new fiber generation will come along every 21 months or so and when it does, we can rapidly and cost-effectively deploy it in the next conduit.

COOK Report: Would you define for me what you mean by a new generation?

Morley: There are many new developments occurring in fiber technology. For example, there is fiber that manages chromatic dispersion.

COOK Report: In other words fiber that inhibits dispersion of the light pulse at high speed?

Morley: Yes. There will be continual improvements in fiber. We announced a couple of months ago a partnership with Corning to be continually developing, researching, and deploying new fiber technologies. The critical thing will be to be able to deploy the fiber and optimize it with new optronics and achieve, by doing so, lower unit costs. And that is what we mean by an upgradeable network. Our conduits will allow us to move down the cost curve multiple times, getting better and better unit costs over time. This is why people are looking to outsource their infrastructure to Level 3. We can pass along this unit cost savings to them over time. They see that they can choose a partner which will move them to lower and lower unit costs over time.

Getting Customers

COOK Report: Have you made a deal with anyone that involves outsourcing over a multi-year period and will cover

one or more conduit upgrades? Or is this something that, for the time being, is somewhat taken on faith?

Morley: The bottom line is that when you are making a decision to work with someone there is so much involved in terms of connecting our networks in making sure our procedures work together, marrying up over provisioning a systems and so on that you are making an enormous investment in a supplier when you do that as a customer. You don't want to make that kind of investment and then two years later have to go look and making it all over again from scratch with a new supplier. you want to be certain that choose supplier will be able to serve you over the long haul. As for whether or not you enter into a long-term contract, you will find that generally customers are looking for long-term partner. These long-term issues are important considerations from the customer standpoint.

We can and do give the potential carrier customers that come to us today pricing that no one else can match, because we are the industry's low-cost provider. Our upgradeable network insures that we will continue to be the industry's low-cost provider. Furthermore with each network upgrade our costs will become increasingly lower than those of our competitors.

COOK Report: I see this as a long-term strategic plan but, on the other hand, opening at second conduit and filling it is not to be cheap proposition. Has anyone looked at the cost and said anything about when you would do this? How close are we to you, or anyone else, actually filling a second conduit?

Morley: Based on technology developments under way we believe that we will be pulling a new fiber probably by the end of next year. And deploying a new generation of optronics on top of that. Once you have a conduit deployed, pulling the fiber is a very very low incremental cost if you compare that to the alternative that the others have which is basically trenching and over-building their networks. That is it's billions of dollars

verses a hundred million dollar cost .

COOK Report: Should one take of 100 million as an approximate cost of your next conduit fill?

Morley: No. It depends on a lot of factors. I can tell you it will be an order of magnitude less than the cost of laying new fiber from scratch.

ONTAP

Let me comment a bit on the provisioning capabilities that I mentioned earlier. We have developed a program that we call ONTAP which stands for On Net Transport Activation Process. We are guaranteeing customers that we will provision on net wavelengths within 30 days of their placing the order and private line services within 10 days of order. We are actually provisioning at intervals that are much less than those guarantees. These intervals are a fraction of what the industry is experiencing today for either wavelengths or optical services. The average time is on the order of 90 to 120 days.

COOK Report: How have you accomplished this?

Morley: By means of really hard work and a lot of focusing. This has been the Holy Grail of the industry. We have been really focused on solving the key issues for out sourcing, transport, high bandwidth and IP services to carrier-type customers.

COOK Report: I'm not hearing much in the way of specifics. Do you consider ONTAP to be just a black box?

Morley: It involves internal process design, internal system development, proprietary system development and interfacing with network elements and configuration and operation procedures. It is definitely a cross functional type of program involving a number of elements across the company. Finally, we are able to successfully implement ONTAP because of our unique highly-scalable and upgradeable network architecture.

COOK Report: If I buy a raw lambda and

want it provisioned as SONET or as gigabit Ethernet and what are my options?

Morley: Typically a customer buys a wavelength because they don't want the SONET protection on top of it. They either want to provision SONET protection themselves or they don't want it at all and will provide IP and protection at the IP layer. Most customers take a wavelength and run it into a SONET protection system. They provision SONET directly over a wavelength using their SONET based ADM equipment. In other cases, they're taking a wavelength and running it straight into a router and they're running IP directly over it.

COOK Report: Running as Ethernet instead of SONET?

Morley: Currently we provision wavelengths with SONET framed interfaces. Today the wavelength would plug into an OC-192 port on a router or into an ADM Transport device.

The Lambda Market

COOK Report: Could we try to conclude then with a summary on your part of what you're lambda's sales look like now, what they are expected to look like a year from now and are they really mainly just to IXCs? Do you have any ILECs or large corporations as customers? What does the market look like?

Morley: The market is still an emerging one. Wavelengths are relatively new to the market and new to customers. I think, in general, customers are still exploring and trying to understand what they can do with wavelengths. However, we are seeing significant interest from carrier-type customers, who we expect to be the initial customers for these types of services, as well as, customers focused primarily on IP services.

COOK Report: What are the terms of purchase? Do I have to sign a contract for a year, or three months or for 6 months?

Morley: You must purchase a minimum one-year term for wavelength. We sell five-year IRUs and nothing under one

year.

COOK Report: Do not foresee selling anything less than a year?

Morley: We do envision the ability for customers to bring this capacity up very rapidly and eventually have a bandwidth on demand kind of capability. This is why our ONTAP system is so important to us. It is a critical step to get to the point of bandwidth on demand.

COOK Report: If I wanted to bring a lambda from 2.5 up to 10 gig, would I find it easily done?

Morley: Depending on how you configured it, it could be basically as easy a matter of changing the interface. A cross connect between here and there and you're ready to go. We provide both portability and upgrade options to our customers. Customers are typically carriers, large ISPs, CLECs and ILECs.

COOK Report: How long ago did you start lambda sales?

Morley: We introduced them in the fall of last year.

COOK Report: You have any figures on what it those sales represent for you now?

Morley: We're not releasing wavelength specific figures but they're definitely a key part of our revenue stream for the year and transport services in total should be about 45% to 50% of our ongoing communications revenue in 2001.

COOK Report: In your quarterly financial reports have you been taking your revenue figures and apportioning them according how much comes from Softswitch services, from managed SONET, from lambdas sales and so on?

Morley: We do not break it out to that level. Revenues are divided between Transport, Softswitch and IP Services. Transport should be about 45%-50% of our ongoing communications revenue, Softswitch about 25%-30% and IP services about 20%-25%.

Cable and Wireless Cuts off Peering with PSI for Five Days -- Both Networks Get Low Marks in Nanog Debate on Peering Policy

Editor's Note: On Monday June 4, the NANOG list discovered that Cable and Wireless had depeered with PSI which a few days previously had declared Chapter 11 bankruptcy. According to **Peter Jansen** of Cable and Wireless (posted to NANOG by Jansen on June 7): "As part of the review of our current peering arrangements, last Saturday (June 2nd), Cable & Wireless disconnected PSINet's peering connections as they no longer meet our publicly stated peering requirements (www.cw.com/us/peering). Over the last few days, discussions have continued between Cable & Wireless and PSINet to find a resolution to this issue for the benefit of all parties. On June 5 PSINet agreed, through a letter of intent, to meet Cable & Wireless' peering policy. PSINet has come forward with new information that may demonstrate PSINet meets Cable & Wireless' peering policy. As an expression of good faith, Cable & Wireless has reestablished the peering connections with PSINet for a period of 60 days while the parties consider PSINet's new information and, if applicable, negotiate in good faith a new peering agreement." Note surprisingly this kicked off several days of lively discussion. We present some high lights below.

June 4, **Moe Allen:** Affective this morning Cable & Wireless started de-peering with PSINET. More when I receive a reply from: peering.admin@cwusa.com

Sean Donelan: Since PSI still hosts one of the root servers (C.PSI.NET) C.ROOT-SERVERS.NET, I hope C&W's customers understand the all ramifications of C&W's actions. Depending on the current state of the net, it may vary from a minor reduction in access to all possible servers to 1/13.

On June 7 **Donelan** added: The loss of one or two (or as someone once calculated up to 40%) of the root servers has no or little effect. DNS is relatively robust,

but not indestructible.

However, no one outside of Cable & Wireless knows just how many networks or even which networks they eventually will cut-off. If C&W continues down the road of "de-peering", they may eventually cut off too much and cease to be a useful Internet provider. They may be a fine private network, but if you eliminate enough interconnectivity you aren't in the Internet any more.

PSI was simply the most public case. But it doesn't appear that even Cable&Wireless's own sales force knows how bad C&W's network connectivity is going to get until after C&W finishes cutting off all the networks. There have been private reports that C&W sent out letters to several other providers.

Will C&W still be a viable network after its management finishes? I don't know, I don't know if they know. They didn't seem to understand the effect cutting PSI would have.

June 4, **Simon Lockhart:** Well, I assumed PSI and/or C&W would have had some backup transit, if they were going to play that game, but...

Donelan: By definition a "tier 1" provider does not have any transit, backup or otherwise. If tier 1 providers terminate peering, there are no alternate routes between them. However, customers who are multi-homed may be able to reach both of them. Likewise customers of other peers can reach both of them. But customers exclusively connected to one or the other battling peer can't reach customers exclusively attached to the other. It's been a while since we've had a real "tier 1" peering battle. Last time ANS lasted a couple of weeks before they caved in and joined the CIX after the CIX filtered out non-member routes.

Vivien M: It can get much more interesting if PSI/Sprint peering was to go down, since at least one provider (whom I won't name, but most people probably can guess who I'm talking about) with some large datacenters (and big customers in those datacenters) uses Sprint to reach PSI, probably ever since PSI attempted to charge the provider with the large datacenters for peering with them. (Another similar large datacenter provider uses Verio to reach PSI, I believe)

I suppose now PSI gets to learn the hard way what happens when they scared half their peers away (to be polite...), and now find that a bunch of the other half are now turning down their PSI peering links. (BTW, has it been established here whether PSI or CW is to blame for this?)

I really really hope that no one still resells PSI dialups, otherwise their tech support lines could be very very busy very soon. Oh, and FYI, a friend also in the UK using PSI reports the same thing you're reporting... no trace of the CW network from there.

Kevin Loch: This is sure to create a number of "big pipe orphans". I wonder if this will create a surge in multihoming attempts?

How could you discourage that now? Unlike the recent DSL disasters, you can't just say "buy a T1 if you want reliable service". Even if you are not a PSI customer, it would be foolish not to multihome now.

Eric Gauthier: I don't know about Abovenet, but when things when down between Exodus and PSI, my impression was that Exodus just got Sprint to carry the traffic. No new circuits, just a new path, and not a big deal because it was small amount of traffic (rumored to be <90Mb).

Vivian M.: AboveNet did the same deal with Verio instead of Sprint... and they're still doing it, too, just like Exodus is still using Sprint to reach PSI.

I don't think the issue here is one of circuits or anything, it's more one of embarrassment. AboveNet, PSI, and probably C&W (I'm not sure about Exodus... ironically enough, it's the only one of these that I use) all claim that they're "tier 1" networks. However, AboveNet has been forced to get Verio to provide transit to PSI because of this. That, technically, means that AboveNet is not a tier 1 by my definition (according to me, and probably most people on this list, a tier 1 is someone who has no transit from anyone). Now, PSI, which used to call itself "the Internet supercarrier" IIRC (ironically, until a year or two, maybe three, ago, also claimed their DS3 frame relay network was state of the art), may be forced to get someone to transit the 2.5 megabits (or is my guess too high?) of traffic to CW. It's not likely to be a big technical deal, but the irony I find to be quite prominent. First PSI forced others to make transit arrangements because of their greed, and now CW is possibly making PSI do the same, for probably the same motives.

Editor: Meanwhile on Friday June 1 some people had noted that there was an apparent peeing problem between Cable and Wireless and PSI.

Sean Donelan: Since on the Internet the sender pays for sent traffic, and the receiver pays for received traffic, I've never understood the argument advanced by BBN/Genuity, UUNET and now apparently C&W that unbalanced traffic means someone is getting a free ride.

If the customer pays flat-rate, you collect the same amount of money no matter how little traffic they send or receive. The 95% charging used by some providers is the greater of *either* inbound or outbound traffic. So imbalanced traffic to or from your customers is paid by your customers.

So, can anyone explain why C&W, UUNET or Genuity care about traffic balance, other than to limit competition

by providers who are better at attracting particular types of customers than them? If you are good at being a webhoster, your traffic will have one profile. If you are good at being an access provider, your traffic will have another profile.

If you are mediocre at everything, I guess your traffic will be balanced.

So, can anyone explain why C&W, UUNET or Genuity care about traffic balance, other than to limit competition by providers who are better at attracting particular types of customers than them? If you are good at being a webhoster, your traffic will have one profile. If you are good at being an access provider, your traffic will have another profile.

Vivian M.: The real reason is probably \$\$\$ plain and simple, but...

My understanding, based on talking to some people who run networks like @Home which are totally access providers, is that the theory they use it this. Let's say you have network A, a big access network, and network H, a hosting network. If the two networks peer in San Jose, Dallas, Chicago, New York, and Washington, DC, and network H's biggest data centers are in San Jose but network A's biggest customer base is in New York, that means that network H sends lots of traffic through the San Jose peering link, and then network A needs to carry tons of traffic on their backbone all the way to New York. Meanwhile, network A sends acks and similar things to network H, and a majority of those go through the New York peering link, and are then taken back to San Jose on network H. The problem, the way network A sees it, is that they might need to get an OC48 between San Jose and New York, whereas network H can get away with an OC3/OC12 on the same path. Thus, network A finds it unjust that they have to pay all this money for this OC48 when network H, which is the network sending them all this traffic, can get away with a much cheaper circuit, and thus they use this excuse to try and bill network H in order to make as much money as possible. Thus the "free ride" argument..

If traffic is balanced, then I guess each

side feels like they're using equivalent amounts of the other's network, and thus that it cancels out... (of course, in a real life implementation, this is presumably significantly more complicated, and I think at that point the logic vanishes in favour of simple greed.)

Mike Leber: You have the cart before the horse (effect before cause), there are really two principles that come before the example policy effect above. They are truisms.

For the purposes of the rules below the term monopolistic peering refers to core networks that have policies that would limit their peers to 10 or so networks IF they were uniformly applied to all their current peers (which they are invariably not (even though a few might be snubbed for general purposes of crassness, ergo C&W depeering a few arbitrarily)).

1) The first rule of monopolistic peering is that the policy MUST overwhelmingly favor the writer of the policy. This is a truism, no company defines a policy that requires them to pay settlements, only vice versa.

For example, this means that for all the lip service paid to settlement based peering compensating the parties equitably you will not find any that allow both parties to be paid. The few settlement based peering contracts I've heard of typically are written as transit fading to settlement free peering when some goal is met (the "correct" ratio and the "correct" quantity of traffic). That is, only one side ever gets compensated. If UUnet offers you a settlement based peering contract you can bet they will not extend you the exact same compensation if you manage to turn the tables on them and attract customers that suck traffic instead of push. Otherwise people would rush out and implement that as a business model.

2) The second rule of monopolistic peering is that the policy MUST be written in a way that allows you to severely limit who you peer with. This is a truism, in order to get the number of qualifying networks down to 10 or so one must write sufficiently restrictive policies.

See, Sean, your error was assuming the rationalization accompanying the rules you see in the more restrictive peering policies legitimately represents the effect of the the policy. heh. :)


Hrmmm...

By the way, many large networks are not monopolistic. Some have hundreds of peers and continue to evolve and grow their networks, taking into account new market entrants that become sufficiently established.

Presumably these networks will provide lower latency and more direct routes (asuming that is what you want) than once upon a time tier one networks that have aggressively shrunk the number of strategic business partners (um, peers, yeah them, who without your network becomes rather boring).

Leo Bicknell: There are several reasons to care about traffic ratio. Where I think the mistake is made is that providers are looking at ratio, but that the ratios they use are fixed regardless of the type of network they are evaluating. That said, it's hard to get more flexible guidelines past the lawyers and bean counters, particularly in a large organization.

Here's a few interesting cases, first, the ratio problem.



Consider "A" is west coast, "B" is east coast. User requests flow B1, B2, A2, while reponses flow A2, A1, B1. Provider 1 ends up carrying more bits a longer distance, and thus incurs a higher cost.

There are several responses to this argument, each with their own problems:

* That's what you get for building and end user network. If you don't like it, build a data center network. Most people don't like suggesting their business

model is broken.

* Use BGP MEDs to make the return route A2, B2, B1. This moves the cost to network 2, which may or may not be fair. Many times provider 1 does not trust provider 2 to do this properly. Even when they do, sometimes it is impossible. BBN and ATT are good examples. If someone sends you a single /8, you have no choice but to hot potato it out, as meds make no sense. The only solution is deaggregation, which has a large number of other problems.

* A settlement should be paid from network 2 to network 1. This is possibly acceptable, if it comes in the form of a settlement. Often the pricing resembles transit, below.

* Network 2 should buy transit from network 1. Most of the medium to large networks are trying to be transit free, and reject this outright. Also, it's quite likely they would by transit from, well, anyone else just so network 1 doesn't get the money from it.

There is an important factor here many of the depeering crowd are missing. The overall traffic ratio of your network is more or less fixed, and is determined by your customer base. Unless you can convert peers to customers (which I have never seen someone be successful in doing on any scale), you will simply move the problem around. That is, if you're 2:1 with Sprint, and depeer 5 10:1 guys, they may well buy transit from Sprint, moving them to 3:1 (due to traffic volume). Now what, depeer Sprint?

Most people from their billing software can add up all customer in and all customer out. If your ratio is under that number, you will NEVER reach it, no matter what you do. Since individual peers will be different, you probably want your limit to be about twice your customer ratio, at a minimum.

This is why I believe you have to evaluate people based on value. Consider someone like @home peering with someone like Globix. One is a pure end user provider, that in fact prevents most

of it's users from running servers and the like. The other is a pure web hosting company, with lots of content and almost no users.

These two networks cannot exist without each other_. If they refused to peer with each other based on ratio, it would be utter folly. Clearly there is great value to both of them in peering, even though the ratio may well be 10:1 or higher.

One of the funniest results of the ratio dance is that it may well create more competitors for a large network. A tight ratio (eg, 1.5:1) is really a requirement that you have a similar customer mix, so you have a similar amount of in+out traffic. How many web hosting networks, who didn't want to compete for end users have been forced to go after end users big time? How many access only networks have done things to attract server users? Companies that could have enjoyed much less competition have forced people to compete with them by ratio.

Equally interesting to me is the "minimum traffic" numbers that many large networks want to put forth. Some of them are quite high, with major networks requiring well over a gig of actual traffic to qualify for peering. This has the effect of pushing the restrictive peering policies down to smaller providers. If a smaller provider has a lot of peers, they send less traffic to any individual peer. One of the easiest ways to get that traffic level is to pull peering with a bunch of transit customers of the network you need to increase traffic with, which of course increases your reliance on that paritcular peer.

One could wonder if some large providers pushed C&W due to the lack of traffic between their networks (since we know C&W had some issues where they couldn't grow their network last year, and had trouble turning up new customers) and that wasn't one of the catalysts for this most recent action.

Vivien M.: I know this was beaten to death a few threads ago, but there are ISPs (eg: one of ours) that bill 95th percentile on the TOTAL of inbound and

outbound.

Albert Meyer: I almost got caught by this one a few months ago. I was fixing to sign a contract with Exodus for a 100bT circuit when I noticed some funny-looking language and asked some probing questions, and then realized that I had to double their quoted rates before comparing them to everyone else. This moved them from the front of the pack back to UU-land. UUNet is another story. They not only charge significantly more than everyone else, but they calculate 95th percentile on the higher of incoming and outgoing rather than the average. When I asked my salesperson why she couldn't give me a competitive rate, she said "Because we're UUNet." She seemed pretty taken aback when I explained to her that UUNet actually had a pretty bad reputation in NetAdmin circles and I wasn't interested in paying a premium for their name. She still declined to give me a competitive rate.

I hear that InFlow charges for average traffic rather than 95th percentile. They're not a backbone, but I wouldn't be surprised to see backbone networks start doing that before too long. It would require some excess capacity, but they would probably make more money in the long run.

Andrew Odlyzko: The discussion over the last few days has been fascinating. I would like to inject a few remarks. The main one is that there is a huge literature on pricing of communication services and other utilities. The problem is that statements such as "Ideally, the price should match as closely as possible the actual cost to provide the service." are very appealing, and even agree with conventional economic doctrine, but founder on the difficulty of determining what "the actual cost to provide the service" is. In addition, there are other conventional economic arguments that say prices should match not costs, but rather willingness to pay.

Electricity pricing, which has been mentioned here several times, is an interesting case. Original pricing was flat rate (so many dollars or cents for each lightbulb

for each month). That was discarded very quickly, largely because of a crucial factor that distinguishes electricity production from telecommunications, namely high marginal costs. Except for hydro (and to some extent nuclear) power, there is a measureable and substantial cost in paying for the fuel that provides each kilowatt-hour of "juice." Even that, though, does not deal with the issues of fixed costs (the generating plant and the transmission lines). How to allocate those costs to various consumers led to extensive discussions and experimentation about a century ago, far more sophisticated than anything that was done in telecommunications at that time or even now. What we have today is usually a combination of fixed rates (dependent on capacity of link) and usage charges (straight fees per kilowatt-hour). Experiments with time-of-day pricing have had mixed outcomes, and there is little of it going on. This might change with smarter appliances, but then it might not.

The Internet does not have the high marginal costs that electricity involves. Hence the economic case is different, but even so, theory does not provide an unambiguous answer as to what the answer is. I have several papers that discuss pricing of Internet and other communications (and even more general) services. As a result of the investigations described in those papers, I do come down on a particular side of the debate (namely in favor of flat rates), but also provide extensive references for other arguments. There is a short, 6-page extended abstract entitled "Internet pricing in light of the history of communication,"

<<http://www.research.att.com/~amo/doc/history.communications1.pdf>>,

(to appear in Proc. ITCOM 2001), the full 40-page paper "Internet pricing and the history of communications,"

<<http://www.research.att.com/~amo/doc/history.communications1b.pdf>>,

to appear in "Computer Networks," and a longer yet and more detailed 160-page manuscript "The history of communications and its implications for the Inter-

net,"

<<http://www.research.att.com/~amo/doc/history.communications0.pdf>>. (Replace .pdf in the URLs above with .ps if you are a fan of PostScript.) These might provide some amusement and possibly even enlightenment.

Security Consultant Chronicles DOS Attack

Thanks to Ed Gerck for calling our attention to The Strange Tale of Denial of Service Attacks against GRC.com. This document found at <http://grc.com/dos/gredos.htm> makes superb reading. An excerpt follows.

It is impossible for an application running under any version of Windows 3.x/95/98/ME or NT to "spoof" its source IP or generate malicious TCP packets such as SYN or ACK floods. As a result, Internet security experts know that non-spoofing Internet attacks are almost certainly being generated by Windows-based PC's. Forging the IP address of an attacking machine (spoofing) is such a trivial thing to do under any of the various UNIX-like operating systems, and it is so effective in hiding the attacking machines, that no hacker would pass up the opportunity if it were available.

It is incredibly fortuitous for the Internet that the massive population of Windows-based machines has never enjoyed this complete "Unix Sockets" support which is so prone to abuse. But the very bad news is . . . This has horribly changed for the worse with the release of Windows 2000 and the pending release of Windows XP.

For no good reason whatsoever, Microsoft has equipped Windows 2000 and XP with the ability FOR ANY APPLICATION to generate incredibly malicious Internet traffic, including spoofed source IP's and SYN-flooding full scale Denial of Service (DoS) attacks! (See my WinXP & DoS Page.) While I was conducting research into the hacker world following these DoS attacks, I encountered evidence — in attack-tool source code — that malicious hackers are already fully aware of the massive malicious power of the new versions of Windows and are waiting impatiently for the "home version" of Windows XP to arrive in the homes of millions of less clueful end users.

Ex-Employee Files \$150 Million Lawsuit Against NSI Alleging Knowledge of Violations of ICANN Agreement and Other Wrong Doing - NSI Closes Domain Policy List, Mueller Critiques ICANN's Position on Root

Editor's Introduction

Those who care enough about the fate of an Internet open to and useable by small business and especially by independent press should continue to be extremely concerned as ICANN's masters continue to propagate the big lie in their attempts to paint what they are doing as 'open' and 'consensus driven,' and a model vehicle of industry self regulation. ICANN is industry freed to doing the bidding of its biggest cronies and to say to the rest of us: to hell with you.

We have to wonder where are the protests of those stalwarts of ISOC and the EFF who in 1997 and 1998 shrieked for the scalp of the evil monopolist NSI saying that ICANN was necessary to bring NSI under control. Three years later we have in place ICANN as a global regulatory authority for the Internet with a multimillion dollar a year budget and staff and agenda that will soon rival that of the ITU.

ICANN gave the appearance of competition by creating other entities that could register names in .com. But all of these so called competitive registrars are but \$6 a year per name fee collectors for the NSI registry which with 30 million names is now guaranteed a whopping \$180 million a year income. Out of this NSI pays a yearly tax to ICANN. Industry self regulation ICANN style means turning the entity you pledged to control on behalf of the internet into your largest single source of funds to enable you both to pursue your continued agendas. The blind deaf and dumb Bush Department of Commerce at the end of May gave final approval to the plan that Joe Sims hatched in January that has enabled NSI to continue in control of its registry and remain a registrar. Vint Cerf and his merry GIPsters will seem do whatever is necessary to keep NSI- Verisign happy.

After all when you have broken as many promises as these men what is one more?

What follows is a slice of life inside NSI which although technically accountable to ICANN in reality is accountable to no one.

Ex-NSI Employee Alleges Illegal and Anticompetitive Actions by NSI

From ICANNWatch we read: **Lawsuit Alleges Illegal Practices in Domain Name Disputes, Allocations** Posted by michael on Tuesday, May 22 @ 11:24:44 MDT Contributed by Michael <<http://www.icannwatch.org/article.php?id=169>>

Someone posting from a hotmail account has sent the domain-policy list a set of documents that include a pro se complaint in an employment discrimination lawsuit filed by Michael Johnson, a former NSI/VeriSign employee. Johnson used to run the UDRP side of NSI's operations, and attended ICANN meetings as an NSI/VeriSign representative. In addition to alleging racially offensive and discriminatory conduct, the complaint includes several paragraphs of spectacular allegations about illegal and preferential conduct by NSI/VeriSign relating to domain name policies and disputes. (See inside for some choice quotes.) Ironically, the domain-policy list is hosted at NSI. UPDATE: By evening, the e-mail from "flip levin" had vanished without a trace. So see inside for full copies of the e-mail and attachments.

NSI counsel Phil Sbarbaro confirmed to me today that Mr. Johnson was a former employee, and that he had filed a lawsuit; otherwise Mr. Sbarbaro would say only that "we never comment on ongoing liti-

gation."

[**Editor's Note:** Sbarbaro undoubtedly is a "good" lawyer. One who will do whatever it takes to win – including intimidating Michael Froomkin into removing the Verisign logo from ICANNwatch's description of the case. See <http://www.icannwatch.org/article.php?id=170> Discussion of the earlier article on ICANN watch brought up a pointer to <http://www.nsi.sex.com> Readers will find posted there 13 pdf files of legal filings made as an "Order to show cause why the court should not bar Philip Sbarbaro, Kevin Golden, and Hanson & Mollowy from further representing NSI. Case No. C 98 20718 JWPVT ENE Kremen vs. Cohen. The papers document Sbarbaro's behavior toward an expert witness hired by Gary Kremen to testify against NSI.]

Here are some selections from the complaint in Johnson v. VeriSign.

50. Plaintiff has personal knowledge of instances of violations of the ICANN Agreement, predatory practices against competitors, disregarding the law when compliance therewith should stop wrongful conduct, ignoring conflicts of interest which produce a competitive advantage for Verisign, and its willingness to proffer untruthfulness instead of the complete truth when the opportunity arises.

51. As the ICANN Representative for Verisign, Plaintiff and Network's General Counsel, often were in conflict with senior management on the Registrar side of Network in urging them not to place the Registrar Accreditation Agreement with ICANN in jeopardy of termination should Verisign's conduct be found to be guilty of violations of criminal misconduct, fraudulent or dishonest behavior, or for acts in which its intentional unfair

dealings directed against its customers, employees and competitors are discovered and proven.

52. Instances arose during Plaintiff's employment where he was required to choose between accepting the threat of his wrongful job termination against public policy and condoning, as more specifically focused in this instant action, violations of federal laws requiring Network and Verisign to fully comply in protecting the protected class employees in its workforce.

55. Among instances where Network disregarded the proper procedures to clear a domain dispute for its advantage and providing different and favorable treatment than that afforded similarly situated customers in violation of Network's rules, a powerful Democratic United States Congressman, Ed Markey sitting on the House Commerce Committee with jurisdiction over both Verisign and Network services, was given special treatment to retrieve his lost domain name that had subsequently been lawfully acquired by Network's paying customer who is a speculator. Congressman Markey, was identified by Verisign's Roger Cochetti, Senior Vice President, Policy, as being "the kind of guy who has a long memory and quite capable of retaliating against those who don't treat him respectfully and he's capable of doing anything to Verisign or us". Defendants used a pretext of a "no contact" notice to seize the domain name even though Verisign was paid current by its customer and knew this was the improper procedure.

56. Other examples where Defendant's Verisign and Network officers instructed its employees to act in ways contrary to the law and spirit of the ICANN agreement, in much the same as issues in the instant litigation, involves not fully responding to the request from the National Telecommunications Information Agency when it expressed written concern about security issues on the server, at a time when the officers of Verisign knew from its customer's and internal technical management team of serious incursions from outside sources because of the system's security weakness, and from the active misconduct of internal

sources engaged in cheating customers by reselling domain names for a personal profit. The internal employees were fired.

57. Verisign employees targeted companies, including one Defendant Wolford called "a copy-cat ratbag registrar" in referring to a major competitor, where lack of timely transfer or refusal to transfer domain names is ordered, or issues relating to the reasonable assurance of payment, pre-registration of multilingual domain names, the pre-payment requirements of the ICANN agreements through Verisign's Partner Program, and the hoarding of domain names are parsed at the risk of liability to the ICANN Agreement and is herein offered to showcase the culture of Defendants and its predisposition for not respecting the line not to cross.

UPDATE: NetSol took down the posting. So here is a new copy of the text of the cover letter and the attachments.

From: flip levin {derrat0123@HOTMAIL.COM} Reply-To: flip levin {derrat0123@HOTMAIL.COM} Date: Tue, 22 May 2001 05:37:16 -0000 To: DOMAIN-POLICY@LISTS.NETSOL.COM Subject: Mike Johnson v Verisign and Network Solutions Lawsuit

Good Morning Jaime:

I just got this from my server at the office. Karen knows the lawyers for Mike Johnson and she got copies of everything. Did you know about this? Did you see where these guys steal from customers and lies to ICANN and NTIA and still get their contract renewed? Something isn't right here.

Hey, notice that ICANN and Department of Commerce and NTIA got this stuff before the final approval from the Congress Commerce Committee. Isn't this dishonest? Did Verisign really buy everybody? Who is still watching the store?

Everybody out here should contact Congress to stop this fast train! This just isn't right. Karen said that Congressmen

Markey and Dingel and several Senators are getting together to stop this deal. We can help by contacting these lawmakers and telling how this affects us! d

Attachments in MS Word format

Attachment 1: Department of Commerce Letter Attachment 2: ICANN Letter Attachment 3: NTIA Letter Attachment 4: Table of Contents for Complaint Attachment 5: Ten Point Version of Lawsuit

[**Editor's Note:** we have read all of Michael Johnson's allegations as stated in his complaint. His recounting of life on the inside of Network Solutions is interesting to say the least. Mark Henderson -Thynne offered his comments on the NSI hosted domain policy list as quoted below. About 36 hours later NSI erased the list which had survived more than five years of bitter criticism of the list owners. Interesting that Michael Johnson's allegations quickly killed it.]

Date: Tue, 22 May 2001 18:44:24 -0400 Reply-To: mark@henderson-thynne.com Sender: Owner-Domain-Policy <owner-domain-policy@LISTS.NETSOL.COM> From: Mark Henderson-Thynne <mark@henderson-thynne.com> Subject: Re: Mike Johnson To: DOMAIN-POLICY@LISTS.NETSOL.COM

As another former NSI employee (I recently left NSI after 3 years in channel development and product management, and some of you may also remember me from NetNames and .tm 3 to 5 years ago as well as the IFWP etc) I can confirm that Michael Johnson is a real person and did work at Network Solutions during the time specified in the lawsuit.

Michael seemed like a nice personable kind of guy and as Director of BAO was responsible for a mostly thankless task: Registrant name changes and disputes.

The events Michael describes surrounding the reorganization of BAO mostly match my recollection but beyond that I do not know if the various accusations around the actions and statements of sen-

ior registrar staff are correct (although the Bruce Chovnick rumors are very well known ;-)

To be honest I do not if there was racism involved or whether that is just Michael's interpretation, but there are 3 areas that could, in my opinion (without knowing all the details), have contributed to him "leaving" NSI:

1) Racism (although I did not see any situations at NSI that I felt were racist) 2) He stuck up for the customer instead of the shareholder 3) He was not liked by senior management (you have to be part of the clique to succeed at NSI) Mark.....

P.S. On a personal note I hope to contribute to domain related forums more now I am no longer an NSI employee. I am currently taking off a couple of months to travel and will then be looking for a new job, I am not sure if I will stay in the domain industry but if anyone knows of any interesting vacancies in the Northern VA or SF Bay area feel free to drop me an email off-list. Thanks.

NSI Kills its Own List

From Icanwatch: VeriSign Pulls Plug on Domain-Policy List Posted by michael on Thursday, May 24 @ 09:26:55 MDT Contributed by michael

Without warning, VeriSign today pulled the plug on the venerable Domain-Policy list. The archives seem to have vanished too. The final message to the list from Tom Newell, announcing immediate stoppage, is reprinted inside.

The move comes as VeriSign is newly secure in its monopoly in .com. And, in the past few days the firm was stung by a series of posts regarding a lawsuit by a former employee. While it's true that there are other places to discuss domain name issues (here for example!), pulling the plug in this way removes archives from the reach of Internet historians. The lack of advance notice also makes it hard for anyone else to volunteer to take over the list at a new host. VeriSign certainly has no obligation to host a mailing list for its critics, although one would have hoped they were big enough to take it with less

petulance. But the decent thing to do would be to leave up the archives, and to make the subscriber list available to a suitable volunteer host if one materializes.

Here is the text of the last message on the Domain-Policy list:

Date: Tue, 24 May 01 11:13:49 -0400
From: Tom Newell {tomn@LISTS.NET-SOL.COM} To: DOMAIN-POLICY@LISTS.NETSOL.COM
Subject: List deactivation

This list will be closed effective immediately. When we started the list many years ago, there were no lists specifically focusing on domain policy issues. Today there are a wide range of public lists that address this topic.

Thank you very much for your participation.

Please refer all queries to Brian O'Shaughnessy, of VeriSign's Corporate Communications department at boshughnessy@verisign.com.

—Tom

Editor's Note: this list was the sole means of open communication between NSI and customers. Now thanks to industry self-regulation this channel is gone. NSI can do whatever it wants because it is not in ICANN's interest to care. Consider the following from ICANN Watch <http://www.icannwatch.org/article.php?id=195&mode=thread&order=0>

Registrar War Developing- NSI joins battle

Posted by jon on Tuesday, June 05 @ 19:47:22 MDT Contributed by Consumernet

After numerous complaints to ICANN staff (chief registrar liaison Dan Halloran) about Registrar.com blocking domain transfers there has been a new development.... Network Solutions has caught wind of it and has now trying the same thing.

It works like this: a domain is up for renewal shortly and the registrant decides to switch registrars. The new registrar verifies the transfer with the admin contact. The old registrar says in order to prevent fraud and unauthorized transfers a second authorization is required. Sounds good so far ... even conscientious.

Then the games begin. Numerous customers from Register.com have complained that the "authorization" produces an error. The customer finds out 7-10 days later the transfer was rejected. Calls, e-mail, and support tickets requesting release of the domain are ignored. Then the deactivation warnings start coming telling the customer to pay up or have the domain deleted.

Today I received my first series of complaints from Network Solutions. It seems they have initiated an e-mail verification system that ... you guessed it ... doesn't work. Here is NSI's response that includes an ad for their services (domain name removed):

Because your intent to Change the Registrar of your domain name, EXAMPLE.COM, to Tucows, Inc. was not confirmed, Network Solutions/Verisign has instructed the Registry not to change the Registrar.

We appreciate your continued business and hope that you will take advantage of all the services Network Solutions/Verisign has to offer. We'll soon be unveiling new services to help you manage your domain name. Be sure to visit our website at <http://www.networksolutions.com> to explore everything you need for value and performance on the Internet.

I have already filed 30 complaints with ICANN (Register.com -37, Registrars.com - 3, and Network Solutions - 3) and a warning for customers was issued at <http://TheNic.com/warning/>

Tucows had previously filed a complaint back in April at <http://www.opensrs.org/archives/dis->

cuss-list/0104/0488.html

The One True Root

Meanwhile at the Stockholm meeting ICANN CEO showed himself to be a true heir of Mike Roberts. Interested in power. Period. Here is Milton Mueller's critique from ICANNWatch <http://www.icannwatch.org/article.php?id=200&mode=thread&order=0>

Are Multiple Roots Outside the Scope of ICANN? Posted by jon on Thursday, June 07 @ 19:15:05 MDT
Contributed by Mueller

In the debate over multiple roots, it is common to see assertions like this: "Alt.roots have no rights inside the ICANN controlled namespace, and ICANN has no rights inside the alt.root namespace. They are mutually exclusive." It would be nice if this were true, but it isn't.

Sooner or later, we will have to stop pretending that multiple DNS roots are "mutually exclusive." If they really were, no one would care about the issue. Everyone could do their own thing, and no one's thing would affect anyone else's.

Alternate roots are not, for the most part, "private" name spaces. They are intended to be substitutes for the ICANN root. They are open to public use. They are almost always designed to be compatible with the existing TLD assignments in the legacy root. This in itself ought to be enough to refute the myth of mutual exclusivity.

However, when conflicting name assignments are made, as in the .biz case, it becomes quite clear that one root must drive out another, or lead to some kind of higher-level compatibility arrangement to avoid widespread conflict. This is because the DNS root is characterized by very strong network effects. If one must choose between a .biz supported by a small root system with say 10,000 users, and a .biz supported by a large root with 50 million users, compatibility concerns will push most users into the one with the widest scope of compatibility.

We can ignore alt.roots now because they

are so much smaller than ICANN's root as to be insignificant. But what if an alternate root becomes as large as or larger than ICANN's? This possibility has been demonstrated by New.net.

Prior to Stockholm, I tried to initiate a DISCUSSION of this issue in DNSO, without presupposing a result. At the VERY LEAST, it would have been nice to come up with a consistent policy regarding TLD assignments that are already being used by other roots.

It should not be that difficult to define criteria to guide ICANN as to when it should avoid making conflicting assignments, and when it should not bother to avoid conflict. E.g., if I have jumped on a TLD name two weeks ago, have no customers, and no real operational capacity, there's no reason ICANN should respect such a claim. On the other hand, it seems self-destructive and destabilizing for ICANN to make TLD assignments that conflict with legitimate, established companies. A company like Ambler's IOD, for example, was using .web and claiming it for many years. New.net has many customers both among ISPs and the public. Conflicting assignments in those cases could be construed as nothing more than a monopoly trying to crush its competitors.

In short, ICANN's relationship to alternate roots is an issue very much like Microsoft's relationship to other operating systems or the old AT&T's interconnection with alternate long distance networks. In fact, AT&T DID say, in the 1960s, that alternative long distance networks, if they were to be allowed at all, had to remain non-connected to the public switched network. It was the interconnection of private microwave networks to the public system (by MCI, no less) in 1978 that ushered in the legal and regulatory revolution that redefined the nature of the telephone system to make it more competitive.

If, as Esther Dyson now claims, ICANN is an "antitrust authority," (and I'll have fun with that assertion in a later post) what are we to make of a competition-enhancing agency that not only refuses to recognize that it is in competition, but actively asserts its right to create conflicts that can snuff out competitors, or lead to extended battles over supremacy that will create problems for users?

In Stockholm, that simple policy issue got mixed up - deliberately by some people but due to ignorance in other cases - with the silly notion that there should be no coordination at the root level.

Whatever one's position on how alternate roots ought to be treated, the real tragedy of Stockholm is that open, bottom-up discussion of this issue was deliberately preempted by the ICANN management. ICANN CEO Stuart Lynn unleashed a PR barrage a few days before the meeting purporting to describe an already existing policy, and the Board and Names Council let him get away with it. The so-called discussion draft was nothing more than another assertion that ICANN is the "authoritative root" and a rather scurrilous set of attacks on the legitimacy of alternate roots; i.e., on ICANN's competitors. And of course, as we learned from the Names Council vote, the business interests that dominate the DNSO did not want to discuss it at all.

It is useful to analyze what the myth of mutual exclusivity accomplishes in this regard. For ICANN, the myth gives it a green light to create conflicts with any alternate roots that might threaten to achieve the critical mass needed to challenge ISPs' adherence to the ICANN root. If our name spaces are "mutually exclusive" after all, it doesn't matter if I come crashing into yours with twenty tons of network effects.

Actually, the "mutual exclusivity" myth would not be so bad if the ICANN supporters would simply be consistent about it. If alternate roots are really outside its scope, then shut up about them! As everyone knows, however, ICANN management and certain Names Councillors recognize that the success of alternate roots really does threaten their monopoly control over DNS administration. If alternatives do achieve widespread use, ICANN will have a tough choice: either fight them and create widespread compatibility problems, or coordinate with them and give up a lot of its ability to impose onerous policies on Internet users, registries and registrars.

This is the corner that ICANN has painted itself into. It must now pretend that other DNS roots are outside its scope and utterly irrelevant to its decisions and processes, while at the same time scurrying about Washington and the IETF lobbying hard to convince everyone that they are destructive and dangerous. Hypocrisy, thy name is spelled I-C-A-N-N!

Vint Cerf's Doublespeak on ICANN Elections

As Decision About the at Large Elections Nears ICANN Tries to Find Any Excuse Maintain Grasp of its Insiders

Editor's Note: Make no mistake about it. ICANN will not allow any more open elections for Directors. The new ICANN budget has no line items for an election in the next fiscal year – millions for regulatory control and none for open democracy. In mid May ICANN's behavior was parodied in a web based cartoon "movie" at <http://www.paradigm.nu/icann/> there Vint Cerf is shown as a two dimensional cartoon like puppet carefully following the Jones Day scripts he is given.

The Stockholm meeting just ended showed how ICANN meetings are generally scripted to be public displays of the justifications that the ICANN Board chair and staff have dreamed up as attempts at justifying their rule. Cerf was caught in the following masterpiece of double speak.

Brett Fausett to the ALSC list June 7: In watching the video of the various [Stockholm] meetings over the last weekend, one exchange stuck with me, and I went back to capture it in more detail. There was a conversation during the General Assembly meeting among Vint Cerf, Esther Dyson, and David Johnson in which Vint made the following observation:

The one question that I had for Esther is, looking at the growth rate of use of the net, the projections that I have are that by 2010, three and a half billion people will have access to the Internet by some means or other. It may be laptops, personal digital assistants, Internet-enabled cell phones, Internet appliances of all kinds. The idea that a particular infrastructure service, kind of like the power system or the road system, needs to have a kind of global election to manage the underlying infrastructure is hard to fully appreciate. We don't do that for the power system and we don't do it for the road system. But it's still really important for the people who are affected by policies to have something to say about it. So, I'm wondering whether the committee has thought about the fact that the

scaling might become just impossible to manage in a literally democratic kind of environment and instead you have to look at other ways of getting public comment on policies, which we do by other than just straight democratic elections.

Fausett: The analogy Vint makes to infrastructure services is, I believe, flawed, and the flaw actually underscores an important point.

In the United States, at least, the power system is a regulated industry, typically overseen by political appointees who, ultimately, are responsible to elected government officials. In some states, members of the public utility commission are directly elected. The "road system" is also overseen by national, state, and local elected government officials. Bond issues submitted directly to the voters often approve or disapprove financing for infrastructure projects.

Jim Dixon: The analogy is indeed deeply flawed. The power system and the roads are both natural monopolies. Furthermore the roads are normally built by public money. The Internet is totally different. It's not a monopoly and it is being built almost entirely by private money. Please, no quibbling about peanuts spent decades ago; the Internet we know today is in most parts of the world almost entirely funded by private capital.

Fausett: Except in California (where power deregulation has been famously unsuccessful), we don't allow the suppliers and distributors of power, and the power utilities who sell to the consumer, to operate without oversight, or even without price controls. We also don't allow road contractors to set their own budgets, set their own quality standards for roads on which public traffic will be carried, and build on-ramps and off-ramps wherever they'd like. These are all regulated by elected government officials.

In making the analogy to two highly regulated industries, Vint really highlights one of the flaws in this ICANN experiment in "industry self-regulation" — it might not be able to scale to allow representation by those affected by its policies.

Dixon: The problem with ICANN is very simply that although it is supposedly an exercise in industry self-regulation, the industry concerned has very little influence over it. ICANN cannot claim any legitimacy; it has no endorsement from industry. On the other hand, the user community has total control over the industry. As those of us in the business understand perfectly well, if you don't please your customers, they can and do just walk away.

Fausett: So what does scale? Our existing governments already have in place effective means of representation. If ICANN can't scale to allow effective, elected representation from the user community, then perhaps it ought to abandon the idea of "self-regulation" altogether and return the decision-making to the governments of the world, where there is already an established infrastructure for elected representation.

Dixon: The number of non-sequiturs here is impressive.

* ICANN isn't industry self-regulation

* There is no reason to believe that any system of elected representation from the user community would represent anyone but those directly involved.

* The decision making is not and does not belong in the hands of the US government - the evolution of the Internet is being driven by industry and its customer base, both of which are global.

* The 'in consultation with other governments'is, shall we say, amazingly condescending.

What scales is what we have now: tens of thousands of competing ISPs, even more businesses and organisations involved in other Internet- related activities, using an incredible number of mechanisms and bodies to coordinate the evolution of the Internet. It's messy, it's very difficult to understand, it's completely unregulated - and it works.

It is hard for me to understand what value any scheme of elected representation would bring to the Internet. It's very very easy to understand, of course, what the value would be to the so-called representatives.

Fausett: Make no mistake, I would prefer to see the ICANN ideal of self-regulation and self-representation succeed, but in a study where "no question is off-limits," you ought to consider the ramifications of failing to find an appropriate representation structure for the user community. If ICANN is incapable of scaling to meet that challenge, one very good option for the ALSC to consider is that the technical coordination functions assumed by ICANN be transferred back to the government.

Einar Stefferud on BWG list: BUT! Wait! There is More! In the US, our road systems are variously controlled at many levels, including local streets and roads, county highways, state highways, and National highways and Interstate Highways. I expect that same is true world wide, more or less.

I also shudder to think of the US National Govt controlling the names of local streets and alleys, or even setting up rules for local naming. The only rules I know of is that names must not be ambiguous for National Highways as a group, or state highways as a group, or County roads as a group, or city names as a group.

There is a Govt agency that compiles the results of local naming of all named areas/regions. It is lodged in NIST, and I have talked with this person several

times in connection with my work with the North American Directory Forum, and separately with the ANSI Committee on X.500/X.400 Name Registration Systems back in the early 1990's.

In this work, I was collaborating with Marshall Rose in developing something we called the "Civil Naming System" for X.500 which would rely on the already established NIST compilation of all named areas in the US, to enable people and entities to provide the upper layers of a name registration tree, more or less like the DNS tree, only with tagged elements like

```
<c=us; st ; city=Huntington-Beach; zip1647-5615; name=Stef>.
```

Of course, this string is a terrible "name" and the whole system failed because Distinguished Name registrations cost \$2000 for perpetual registration. One big flaw was the idea of perpetual registrations, which requires the use of the "cemetery lot" business model;-)... Pay once, and the interest on the invested fee then pays to water and mow the grass for eternity.

I think the whole ANSI thing was dropped after it only registered a few thousand enterprise names over several years. It was a huge hassle to get a name registered, costing much more over all than the \$2000 fee.

ANSI was spending more than they were taking in;-)...

The value of all this experience for me was/is to shun all instances of any kind of repeat performance;-)... Fortunately, the US Govt only records and documents what locals do, and does not tell them what to do.

This is where the DNS is going to end up in due course, with lots of messing around along the way.

Attorney BWG: I have doubts that ICANN is going to fall of its own failures, so in the interest of tempering some of its worst practices, I think an elected At Large makes ICANN better, if not perfect. I also think you're mistaken on

the personal value of being a Director to those who are elected; with the possible exception of Jonathan Cohen, I suspect there's more sacrifice than benefit.

Dixon: The ICANN we know is or has been run by Esther Dyson, Joe Sims, Mike Roberts, and so forth. I haven't seen much sacrifice on their part. They do seem to have accumulated considerable benefit.

You seem to be advocating a strengthened ICANN run by representatives elected in a global election. Unfortunately, from what I have seen, those representatives are likely to represent vanishingly small cliques - while claiming to represent huge groups. 99.999% of those huge groups will have never heard of these self-appointed representatives.

Attorney: On the idea that governments are preferable to self-regulation, I don't think they are. They're far worse, in fact. But it's an old Jedi defense litigation trick to make sure every opponent has a downside to pursuing you, and I think the powers-that-be in ICANN believe that abandoning the At Large is only a momentary PR hit, not a complete mission failure that undermines its ability to continue operations.

If ICANN fails, then 'governance' of the root reverts to the U.S. Government, not the UN.

Dixon: Do you really believe that the government of the United Kingdom would just bow down if the US government attempted to exercise any sort of control over .UK? What do you think the French reaction would be to a claim of US sovereignty over .FR?

In the real world, what matters is where the ISPs get their root zones from.

In that same real world, the US government is not going to attempt something so Quixotic as trying to actually alter the way in which the root zone is managed. If you step back and look, the root servers are still where they were five years ago.

Executive Summary:

AT&T, pp. 1 -11

We interview John Strand, and Tom Afferton of AT&T. Afferton directs AT&T's Advanced Transport Technology and Architecture Planning. Strand is a consultant for the Optical Networks Research Department at AT&T. He also is a Working group Chair in the Optical Internetworking Forum. The interview describes the choices faced by AT&T in constructing an intelligent optical transport network to take advantage of its new backbone overbuild and the developments of the past five years in optical transport technology.

The amount of fiber available to carriers has exploded as the result of builds carried out largely by new players since 1996. As John Strand points out, by 1995 the optical transport networks of AT&T, Sprint, and MCI accounted for 75% of total inter city fiber deployment in the US. By the year 2000, MCI, Sprint and AT&T fiber had shrunk to less than 1/3 of the national total with the remainder belonging to no less than 39 new national carriers. And, as the year 2000 began, there was an estimated total of 400,000 route miles of fiber with an average 46 strands of fiber per cable.

One may argue that one possible reason for this build out and the consequent huge amount of debt taken on by the telecommunications industry is a misunderstanding of the growth rates of internet traffic. We cite in detail the argument of Andrew Odlyzko that although for a time in 1995 and 1996 internet traffic likely did double every 90 to 120 days by 1997 growth had slowed down to about 100% a year. The consequences of a difference in growth of between 800 to 1500 percent a year and 100% per year are now being felt both by the owners of the transport and builders of the optronics and routers. See <http://www.research.att.com/~amo/doc/recent.html>.

Odlyzko presents a well-supported series

of arguments as to why the doubling of growth three or even four times a year was a myth. It was nevertheless a myth that shaped a stampede that led to the laying of vast amounts of fiber by companies seeking to win what was an impossible to maintain sprint of growth and obtain a first mover advantage where those who lagged behind would never be able to catch up to the one or two players that quickly became dominant.

Optical technologies did yield 100 times and more recently probably even 1000 times advantages in the cost of moving bits in an abstract sense when compared to the legacy networks of ATT, Sprint, WorldCom and the ILECs. A first mover advantage there could overturn global telecommunications. There were, however, only two problems. First the amount of capital required to rebuild the infrastructure of the global PSTN was simply too vast for any single company no matter how ambitious its borrowing and spending plans. Second, as the largest carriers moved forward to try to win the technology transition race, they moved to buy single-vendor optronics solutions which were the only thing that could be quickly rushed into service.

The development of optical internetworking standards was too often overlooked in the drive to buy whatever solution seemed to fulfill the promise of capturing first mover advantage. Industry players did found the Optical Internetworking Forum in 1998. As we learn from our interview with Strand and Afferton, the work of the OIF will play the critical role in deciding how soon the new fiber networks will be able to really inter-operate and, in like manner, how soon they can rely on an open national market that uses their infrastructure to sell bandwidth rather than rely on what ever small part of national or global demand they can meet on their own.

Meanwhile, AT&T, like many other traditional carriers is faced, with the investment of very large sums of money necessary to make the new optical technology a part of its infrastructure and prepare for the day when the voice traffic that accounts for only about half the total but

brings in about 6 to 7 times as much money as data traffic shrinks to very small minority of total traffic. The problem is of course made more difficult by the fact what is a cost efficient architecture for circuit switched voice is not one for data.

As gigabit and ten gigabit Ethernet begins to blur the difference between LAN and WAN, enterprise customers increasingly will want bandwidth that runs seamlessly from an office in New York to an office in California. One of the major concerns for AT&T, which is rich both in metro area fiber and inter city fiber, is to architect a network where the differing requirements of long haul and metro optronics will not be a barrier to the rapid provisioning of customer circuits. AT&T and the other carriers are well aware of the concerns expressed by Avi Freedman in the May *COOK Report* that provisioning such circuits now takes months.

According to Afferton, a key feature of what AT&T is building relies on "SONET multiplexers that have added intelligence and can handle not only SONET TDM type traffic but also, for example, can map Ethernet into SONET, and can do some ATM switching and so forth. We have started rolling these out into our metro networks in order to reduce our costs, as far as SONET traffic is concerned, but also to have the intelligence to inter-operate with the intelligent optical switches so that we can then do this automated provisioning end-to-end."

Rapid provisioning and ultimately giving the customer the opportunity to do the provisioning is the goal by year's end. The current economics of the technology transition dictate that any carrier with a large amount of fiber will be selling what ever flavor of bits including lambdas it can find customers for. AT&T has been selling lambdas since 1999 as it continues to upgrade its network with switches from Lucent and NEC.

Intro to Level 3, pp. 12 - 17

We interview Andrew Morley, Senior Vice President of Global Strategy at

Level 3. Morley recounts Level 3's impressive global infrastructure and lines of business. One model for survival in the industry's current debt crunch is for a company to build a global green field infrastructure so substantial and cost effective that it can convince the traditional carriers that it is more cost effective to outsource new data and perhaps even voice services to Level 3 than to invest the huge amounts of money necessary to upgrade their own infrastructure further.

Based on additional interview that we have done with Ron Vidal and Rob Hagens at Level 3 there is evidence that Level 3's preferred way out of the present situation is to adopt a kind of horizontal business model as infrastructure supplier to legacy vertically integrated carriers. In any case it needs to gain as much income from its investments as it possibly can so that it can pay the debt on its bonds before it runs out of cash. Given its resources if it cannot do this, the portent for the future of the industry is especially grim for Level 3 would currently seem to have more going for it than the rest of the new green field players.

Still the need for increased cash flow has resulted in tightened budgets and in a degree of ambiguity about its mission. According to Morley, Level 3 has 3 major lines of business: transport (sale of lambdas and dark fiber), Internet Protocol services and Softswitch services. For 2001 he declared that "Transport should be about 45%-50% of our ongoing communications revenue, Softswitch about 25%-30% and IP services about 20%-25%."

Softswitch is one of L3's biggest success stories with sales of its managed modem Softswitch services running at 8 billion minutes per month. A different source informed us that these services in the past year have for AOL and EarthLink brought the monthly cost of maintaining a dial up port for internet service down from 15 to 3 dollars. To keep its large enterprise customers happy Like AT&T, Level 3 is focusing on rapid provisioning of bandwidth with a claim that its guarantee of less than 30 days from order to turn up of service is the best in the industry.

try.

PSI versus C&W Peering Spat, pp. 18- 21

From June 2 through June 5th Cable and Wireless turned off peering with now bankrupt PSI. We publish an edited version of some of the discussion on NANOG which shows that the reputations of both parties were tarnished. Peering is still a matter handled privately and without any concrete standards that offer much assistance to those who want it. Because it is a vehicle without firm rules, most service providers feel that it discriminates against smaller players and leaves the industry prone to significant disruptions in service. peering disputes are one of the few occasions when players will 'vent' enough to give the rest of us an idea how these critical relationships affect the rest of the industry.

Excellent Description of DOS Attacks, p. 21

Steve Gibson, Gibson Research Corporation has published on his website: The Strange Tale of the Denial of Service Attacks Against GRC.com <<http://grc.com/dos/grcdos.htm>> "Nothing more than the whim of a 13-year old hacker is required to knock any user, site, or server right off the Internet.

I believe you will be as fascinated and concerned as I am by the findings of my post-attack forensic analysis, and the results of my subsequent infiltration into the networks and technologies being used by some of the Internet's most active hackers." Gibson ends his essay: We need a tool to hold ISP's accountable and publicly demonstrate individual ISP responsibility. Given the universal reluctance they have demonstrated so far, I believe that only active public scrutiny will bring about the changes required to insure a reliable and secure future for the Internet. The development of that tool is my next project. The name of that FREE tool will be: SpooFarino tm"

ICANN-NSI Topics, pp. 22 - 28

NSI hit with \$150 million lawsuit. We summarize the events of the release of Michael

Johnson's allegations against NSI ranging from "not fully responding to the request from the National Telecommunications Information Agency when it expressed written concern about security issues on the server, at a time when the officers of Verisign knew from its customer's and internal technical management team of serious incursions from outside sources because of the system's security weakness" to "personal knowledge of instances of violations of the ICANN Agreement, predatory practices against competitors, disregarding the law when compliance therewith should stop wrongful conduct, ignoring conflicts of interest which produce a competitive advantage for Verisign." ICANN was supposed to ride herd on NSI. Instead in the absence of any oversight ICANN stands by and collects its tithe from NSI which, having just been given an essential permanent monopoly over dot com without warning shut down the mail list on which the Johnson allegations were published. We would like to move cookreport.com to a competing registrar, but having seen the trouble that NSI's system has created for others (including loss of domain) we would not dare attempt it without the assistance of a consultant. Some competition.

We include several pieces from ICANN watch. Milton Mueller's critique Are Multiple Roots Outside the Scope of ICANN? <http://www.icannwatch.org/article.php?sid=200&mode=thread&order=0> Also Registrar War Developing- NSI joins battle <http://www.icannwatch.org/article.php?sid=195&mode=thread&order=0>

Make no mistake about it. ICANN will not allow any more open elections for Directors. The new ICANN budget has no line items for an election in the next fiscal year - millions for regulatory control and none for open democracy. In mid May ICANN's behavior was parodied in a web based cartoon "movie" at <http://www.paradigm.nu/icann/> there Vint Cerf is shown as a two dimensional cartoon like puppet carefully following the Jones Day scripts he is given.

During the Stockholm meeting just ended Cerf said: "The idea that a particular infrastructure service, kind of like the power system or the road system, needs to have a kind of global election to manage the underlying infrastructure is hard to fully appreciate." We publish a discussion that shreds Cerf's labored reasoning.

ICANN Cannot, Say Critics

By Steve Kettmann (from Wired)

June 9, 2001 PDT

BERLIN -- Last weekend's four-day quarterly meeting of the Internet Corporation for Assigned Names and Numbers was such a fiasco, it has strengthened the calls for major structural changes in the board.

Andy Mueller-Maguhn, one of five at-large board members elected late last year, abstained from a vote to approve the minutes of the previous meeting, explaining he had not read any such minutes.

"Fifteen minutes later it was discovered that there were no minutes," Muller-Maguhn said. "But that is the mentality on the ICANN board: Always to say yes. It's a lot like in the old East Germany." The embarrassing phantom-minutes vote -- and the rubber-stamp

image it helps create -- was just one detail that has emerged of a meeting many critics say illustrates the ICANN board's inefficiency and lack of responsiveness to the world's Internet users.

Only five of the 19 board members were elected, leaving a board that many see as skewed in its worldview.

"ICANN's founding premise, as defined by the attorneys who put it together, was that the only people who should have a voice in ICANN were stakeholders, which is essentially a code word that means someone who makes money from the Internet," said Karl Auerbach of California, another elected board member.

"If you're an Internet user, and you end up paying the bills, ICANN's perspective is that you don't count," he said. "It has made the end user the victim."

Editor's Note: For the remainder of the article see:

<http://www.wired.com/news/politics/0,1283,44404,00.html>

Light IP and Gigabit Ethernet - \$375 Now available.

<http://cookreport.com/light-ippgige.shtml>

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**