



Caimis Launched as CAIDA Commercial Spin Off Goal Is to Give Internet Industry its First Production Quality Network Management, Traffic Analysis and Location Sensitive Measurement Tools

Editor's Note: We interviewed Tracie Monk and Kim Claffy (**kc**) more than two years ago not long after CAIDA's launch. With the launch of Caimis it seems likely that for the first time there will be more economically rewarding career paths in the development of these traffic monitoring tools. Even more important though is the creation of a business model that makes production quality tools available for the first time. We conducted the interview that follows on May 17th, well before the launch date of Caimis was known.

Tracie Monk: Let me introduce the people that we have here. We've got **kc** who you know and has been here for a long time, who is the head of CAIDA at the University of California, San Diego's San Diego Supercomputer Center. I am here wearing two hats right now: I am still director of CAIDA at this stage, but I'm leaving the university in June to start up a company by the name of CAIMIS.

COOK Report: CAIMIS Stands for?

Monk: The name actually came from the concept of CAIDA's Monitoring and Internet Solutions. However, the corporate name is only Caimis, Inc. CAIDA's Monitoring and Internet Solutions is what we were thinking of when we came up with the name. We have two other people here. We have David Moore, who is the senior technical manager at CAIDA. And we have Daniel McRobb, who is the author of several of CAIDA's tools, but who will be joining Caimis next week to become the chief architect for the monitoring and network management systems that Caimis will be developing. Daniel will be located in Caimis's Ann Arbor, Michigan office.

COOK Report: Does this have anything to do with Merit?

Monk: No, nothing to do with Merit. The

relationship, however, is with engineers who were a part of the NSFnet backbone who came from ANS Communications. And we have several of the premiere engineers and software developers that will be joining Caimis in Ann Arbor.

COOK Report: And, the engineers who are joining you have been freelancing or working for other people in the meantime?

Daniel McRobb: They've been working for UUnet, who bought ANS from AOL.

Monk: I know that you've heard of some of the tools that Daniel has developed. He is the author of cflowd, skitter, and arts++. See for example <http://www.caida.org/tools/measurement/skitter/>.

COOK Report: Since this is assumed to be an introduction for people who are probably not familiar with this, let it be said that they are tools involved in network measurement, analysis traffic, and so on. For example, according to your web site <http://www.caida.org/tools/measurement/cflowd/> cflowd is a flow analysis tool currently used for analyzing Cisco's NetFlow enabled switching method. The current release includes the collections, storage, and basic analysis modules for cflowd and for arts++ libraries. This analysis package permits data collection and analysis by ISPs and network engineers in support of capacity planning, trends analysis, and characterization of workloads in a network service provider environment. What is your objective then with CAIDA and now Caimis?

Moore: What CAIDA has done in the past is develop prototype measurement devices, tools, and analysis techniques. The problem is that we don't have the staff or the resources or the environment at the university to turn these into actual products that are useful to the commercial sector. The university model basically is that we get

Volume VIII, No. 6, September 2000
ISSN 1071 - 6327

these to a prototype state, they meet our needs, and the university then licenses these out to a company which would go on to turn them into real products. The new company, in this case, Caimis will market and sell them to industry.

kc: The government calls it 'technology transfer'. Getting to the stage of launching Caimis means their funding model worked. The NSFnet had similar effect on a scale several orders of magnitude larger. We believe that Caimis helps the whole community because more ISPs and users will be willing and able to use tools that have real support behind them.

COOK Report: I understand. May I assume that part of the ideas for what you're doing has probably come from feedback from ISPs that you're working on?

Moore: Absolutely correct. There has been a lot of feedback that our prototype tools just aren't ready for deployment in a backbone. They're not documented well enough, they're not cleaned up enough. They're prototypes basically, and we're not a support organization so we can't help people.

kc: So, the target audience is still ISPs and

On the Inside:

Caimis & CAIDA	pp 1 - 8
Cogent Communicatons	pp. 9 - 14
Icann Footnotes	pp. 15 - 21
Optical BGP	pp. 22 - 24
Executive Summary	pp. 8, 25 26

the Internet community, in that we recognize the need for tools that can gather data at a macroscopic level across different ISPs. The spinoff of Caimis is designed to be responsive to what ISPs have articulated as their needs. They tell us: "We'd like to get you [the research community] data, but we can't really put these research tools in our network; they're just not sturdy enough for us." So it is really a natural evolution.

Part One: CAIDA

COOK Report: CAIDA will continue to do research?

kc: CAIDA has several ongoing activities, including research and prototype tools but also activities in education and outreach, membership support, and the SD-NAP access point at SDSC, which has about 25 participant providers. Information about all is available from the main CAIDA web page at <http://www.caida.org/>.

Most recently CAIDA sponsored the 2nd annual IEC (Internet Education Curriculum) workshop, and phase 2 of the IEC-sponsored ITL (Internet Teaching Laboratories). The IEC project is a repository web site <http://www.caida.org/outreach/iec/> where we keep a bunch of teaching and reference materials for classes on Internet engineering.

Monk: The Internet Teaching Laboratory (ITL) project involves donations of equipment from Cisco Systems and Cable and Wireless to create router labs and internet teaching labs at universities throughout the United States. MCI Worldcom is donating engineering support to help make this possible and the NSF is supporting us through funding to do workshops and travel and setup the initial laboratories. We have a fairly impressive board of advisors, including Robert Aiken, Cisco; Steve Bellovin, ATT Research; Scott Bradner, Harvard; Randy Bush, Verio; John Connolly, Univ. Kentucky; Jon Crowcroft, Univ. College London; Rick Wilder, Broadband Office; and Jim Kurose, U Massachusetts.

COOK Report: I remember hearing something about this about a year ago. Is it really just now getting setting up? Are any of the programs in operation at this point?

Monk: There were significant delays in getting the Cisco equipment out of the pops that they were being used in. Everything has now been shipped. The first solicitation was last fall, with 38 proposal submissions from universities, twelve of which were selected to be recipients of laboratories. We are in the process of setting those up. I know that Humboldt is up and running now. I'm not

sure which other ones are fully operational.

COOK Report: And those are on one of the URLs?

Monk: Yes details are at <http://www.caida.org/outreach/itl/>. We are almost done with the second round solicitation that closes at the end of May, and we envision that 13 additional universities will be chosen as recipients. CAIDA will be spending the summer setting up those labs. The ones now funded and in the process of being set up are: The University of California, San Diego; Grand Valley State University (Allendale, Michigan); Humboldt State University (Arcata, California); Kent State University (Kent, Ohio); Syracuse University (Syracuse, New York); The University of Alabama, Birmingham; The University of California, Los Angeles; The University of New Hampshire (Durham, New Hampshire); The University of Southern California (Los Angeles, California); The University of Texas, Austin; Ohio University (Athens, Ohio); and Florida State University (Tallahassee, Florida).

kc: ITL is a sub-piece, addendum to the Internet Engineering Curriculum project, and the inspiration for that one again was derived from needing better curriculum materials in the universities-level right now. And, again all responsive to the ISPs articulating the need to be able to hire engineers right out of school that were useful without years of training. The problem that has existed is that a lot of the Internet curriculum out there is weak or non-existent. It is also not up to date.

COOK Report: And I had the impression that at least a year ago, and I bet it hasn't changed. The goal has been to have a reasonable set of routers that students in the university could play with. It just hasn't been there. Has it?

kc: That's correct.

Moore: Basically the routers are in a closet. They are wired to each other directly and the university has to sign an agreement that they will not use the equipment for their operational networks. They can only be used for research purposes and for class purposes.

COOK Report: I understand that. Am I correct in assuming that one would have to have a couple of leased lines to attach to the routers?

kc: No. That's the university's problem. You don't have to use leased lines anyway. The routers and other equipment are hooked together to form a lab in a classroom. You don't need it attached to anything to be able to configure BGP stuff on it.

David: Basically, they are tabled back to back. So, they are actually setting up networks inside the labs, but there aren't any leased lines involved.

San Diego NAP

Monk: The San Diego NAP (SDNAP) is another project that CAIDA is working on right now. The NAP is here in the building (SDSC). There's currently 25 participating local ISPs that are all peering here. Tachyon is bringing in a lot of customers with their satellite network. And, we work with these companies to get them operational. We do training workshops on BGP and other things to get them up to speed to learn how to actually inter-operate with other ISPs. The condition for being able to come into the NAP is that the traffic that passes through there is available to CAIDA for monitoring purposes for traffic analysis.

COOK Report: Are there other commercial NAPs in San Diego? What would determine whether or not an ISP would want to come in? If you're the only game in town, the ISP would obviously want to come in.

Monk: There are a number of co-location facilities currently. However, there are no neutral 3rd party network access points in San Diego, besides the San Diego NAP here at the Supercomputer Center.

COOK Report: From the point of view of the commercial ISP, does anyone ever raise an issue about making the traffic available for research purposes?

Monk: If they make it an issue, then they don't to us, because it is very pointed on our website that that is a pre-condition of participation. So, if they are deterred by that pre-condition, they're not even talking to us.

COOK Report: Can one assume presumably there are no privacy issues or anything like that involved? I mean, a packet is a packet is a packet?

Monk: No, we do sign non-disclosure agreements with all participants of the NAP. They technically own all of their data. CAIDA is allowed to look at it for research purposes only. We do not publish any of the information. It is all tightly-held information. Details about an ISP's traffic are not made public in any form or fashion.

OC 48 Mon

kc: Let's go back to our research efforts. Basically all, involve analysis of Internet data and development of advanced measurement tools, like OC48 monitor which is on our plate to do for this year. We get requests

from ISPs right now who are putting in OC48 links and routers that they want to be able to monitor.

COOK Report: Supposing you have a 155 link, maybe it makes no sense to monitor that?

kc: You mean OC3 link? We have OC3 monitors right now that can handle both ATM and Packet Over Sonet (POS). Actually, we have OC3 and OC12 support through our Coral monitors. This is a very good example of a research tool that ISPs are a little bit nervous about deploying operationally. Because they're not real products. It's definitely research code. Caimis will be licensing the OC3 and OC12 monitors (Coral Reef code), but since the OC48 monitor doesn't exist yet, there hasn't been any licensing discussions regarding it.

COOK Report: At some point, 6 months, a year, 18 months into the future, do you ever envision, what would it be 4 steps higher, an OC192 mon or is it unlikely to be needed here?

kc: It will be needed with O192 links. It's just not on the plate yet.

Moore: More potentially interesting things at the moment are gigabyte Ethernet and 10-gigabyte Ethernet.

COOK Report: Absolutely. What's up with those?

Moore: We have funding to develop a gigabyte Ethernet monitor. We've been talking to some other groups who are interesting in developing a ten-gigabyte Ethernet monitoring solutions as well.

The Tools and Their Uses

COOK Report: Consider me to be someone who really knows nothing about the area that you are talking about when you say a gigabyte or a ten-gigabyte monitor. Obviously, a monitor is something that looks at a link in some way. But what does it do? It tells the observer how the link is performing? Can you go into a little bit of detail?

Moore: In this case, what we tend to mean by monitor, for our definition, is something that is capable of basically sniffing all of the traffic off that link, or putting in some sort of filtering and saying that I only want to see this kind of traffic. Being able to look at IP headers, potentially even payload for security applications. It's a passive monitoring technique. You tap into the network somehow and you see all the traffic that passes by it and the monitor box can do whatever analysis it needs to do with that data.

COOK Report: What would be the reasons that an ISP would use a monitor box like this? Would they use it for some sort of traffic engineering, some load analysis, or something?

Moore: You could use it for a variety of things. You could use it for workload characterization, knowing what the mix of traffic is, or application break-down. Is there a shift among my users to more and more multimedia applications, and if that's the case how do I engineer my network to meet the needs of multimedia applications?

kc: One of the data sets that ISPs really want is traffic flow matrixes. Where is traffic going to and from so that they can do better engineering capacity planning kind of things. NetFlow export gives you information from which you can derive a traffic matrix. One of the tools that Daniel **McRobb** developed while he was at ANS and CAIDA was cflowd, which basically takes this NetFlow export stream of data out of Cisco routers and turns it into lots of useful objects including this traffic matrix. So that is one thing that ISPs definitely find useful. They have asked for continued support for that in Cisco's and other vendors' products. Cisco's turns out to be the only routers currently that give you anything close to that information and NetFlow doesn't tend to work that well with their higher end routers, like the GSRs. Consequently, if you have a bunch of OC12 links in your network and you want to get traffic matrix information out of them, right now your only option is to get one of these OC12 monitors, a Coral box.

COOK Report: Is one possible idea then that as you go on and develop this, Cisco or others would eventually acquire and use what you upgraded?

Moore: Actually not. The ideal solution is to be able to do as much on the router as possible and not to be forced to put a second box in your POP. But, one of the points that I wanted to make is that the cflowd that **kc** just mentioned is a GPL'd tool. Daniel **McRobb** plans to continue work directly on enhancements to this tool and to ensure that it operates not just with Cisco routers, but so that other router vendor products, particularly Juniper's.

COOK Report: I think you used the term, GPL'd tool?

Moore: That has to do with the copyright. It's GNU Public License. One of the goals of Caimis is to leverage on open-source products wherever feasible and to the extent that we can keep certain critical products in the public sector and available to the community, we want to do so. And, cflowd is one of those.

COOK Report: What would be some of the other products or tools that would fit into that category?

Moore: Another specific set of code that we'll be doing the same thing with is arts++, which is a binary file format specification that is used to store data from a number of our tools. Caimis will continue to enhance this source code and make it publicly available through CAIDA. However at the same time, we will be developing proprietary modules that will sit on top of these public tools.

COOK Report: Tell me a little bit more about how those will work, and the purpose of them, and what's motivating them.

McRobb: Currently, there aren't really any graphical plotting utilities and report generators in the cflowd or arts++ packages. One of the reasons for that is that CAIDA doesn't really have the resources to develop those tools. In addition, those tools tend to require a lot of customization for a particular provider's application, so that's where I expect Caimis to do some work on the reporting and presentation tools side.

COOK Report: For illustrative purposes, can you describe any of the actual or theoretical application that provider X, Y, Z might want to go to Caimis for on this sort of thing?

McRobb: A good example is a billing application. You can use this kind of information for accounting and billing, and of course CAIDA really has no interest in doing that, but the providers do. So, that's an example of an application that Caimis could provide.

Importance of Defining and Understanding the Internet Core

kc: Let me describe some different kinds of research topics that are relevant in the next year. CAIDA recently created a poster to depict a sample snapshot of the 'core of the Internet'. So, for example, how many of the 6,000 autonomous systems AS numbers that are out there are really "core" — in other words, are very highly connected. [See http://www.caida.org/analysis/topology/as_core_network/.] Strongly connected means having a high "out-degree" if you draw a graph. We've drawn large graphs of autonomous systems as depicted from both skitter records and other researchers' data involving many thousands of trace routes and IP paths.

COOK Report: How does this impinge on the technical arguments that one has been seeing about CIDR and distribution of IP numbers and can we afford to let the core

continue to grow larger and larger. Is there an immediate direct relationship to these kinds of discussions?

kc: No, not really. Those discussions are rather diffuse right now I think. We're at a more basic level. Nobody has really even made a definition of what it means to be "core." In other words what the Internet "core" really is.

COOK Report: Help me understand more why this is something that is interesting to do and when you understand it better, where do you think it leads?

Moore: Let's say that you have a lot of topology information from skitter or from routing tables you want to understand the nature of the Internet in some sense. Either just to understand or to help design simulators that model the Internet more reasonably.

kc: Or, to find critical areas of connectivity, like ASs that are responsible for so much of the connectivity with other ASs that if you took them out it would really rip out a whole piece of the Internet.

Moore: There are a lot of leaf-nodes. Universities, even local ISPs, are kind of near the edge of the network in some way and they're not really important in some sense. We're trying to develop techniques so if you say I've got this large graph of inter-connectivity between peering sessions, relationships between ISPs you can find out what ones tend to matter more than the others?

kc: So, if one peering session or one AS or one piece of the network would be removed, how would the ISP or a particular AS's connectivity to the rest of the Internet be affected?

COOK Report: The idea is that if you know something about this then you can do some sort of contingency planning and develop your routing so that if one route goes away you fall back into another connection?

kc: Both operationally to develop more redundancy, if you see what the Internet looks like now, and you want to protect yourself against vulnerable infrastructure. However we also do this for research involving the use of these topology maps as more realistic input into simulators and models.

COOK Report: You talked about protecting yourself from vulnerabilities? Is it at all related to understanding how you might fall back if you had a fiber cut on a major link?

kc: Yes, that's what I meant as an example of ripping out a peering session or an AS link or an AS. Another important area for

research is to correlate skitter forward path data with what is articulated by BGP routing tables. There turns out to be a lot of disparity (some of it expected) between the picture painted by BGP and a router in terms of the AS paths that are expected to be taken through that router and the paths that are actually taken.

COOK Report: And, what is there that accounts for the difference, or is that the whole point of what you're getting at?

kc: There are a few obvious things, but we haven't finished making a taxonomy of all of the reasons and categorizing what disparity is attributable to what reason.

COOK Report: I wonder if this would be relevant to the situation where we to back off and try to get to a 50,000 or 100,000 foot view of the Internet for some investment person who is looking at the general ability of the Internet to continue to scale.

I was talking to, Noel Chiappa, the other day in getting some of his ideas at what would be necessary for the network to be able to do to accommodate a commodities market in bandwidth. This would be a world where in a couple of years from now I might be able to order an OC3 from San Diego to Chicago to be activated at midnight on Sunday or Monday morning and I want the thing to run for 24 hours. One of the comments that he had was that if you hypothesize something like that and you create a world where a lot of people may be wanting to turn off and on these big links fairly rapidly all over the place, that he seemed to have some idea that the routing system we currently have would find it difficult to accommodate that.

kc: It's not a matter of difficult. It's a non-starter. It's more than difficult with the current system.

COOK Report: In other words, it just won't happen. And, I know from talking to Stan Hanks and some other people, there are a lot of people really interested in making a commodities market in bandwidth possible, so what are some of the things needed to make it feasible? It sounds like some of the things necessary are tools that you are developing?

kc: Well, I tend to agree with Noel there, I don't think the current BGP system is equipped to deal with that. I think it really needs to be radically changed from the bottom up, and that's not going to happen very easily because you have to get too many people on board at one time.

COOK Report: To what extent is BGP or improving routing protocols anywhere on your agenda?

kc: It is not. We don't write routing protocols. We're not going to change BGP. That's way beyond our scope. We're still trying to understand the dynamics of the routing behavior, but the protocol we accept as a black box. We're not trying to tweak it.

COOK Report: We've got people out there that have routing problems. We want to do certain things. We've got BGP to do it with. And, if you said in turning all this big bandwidth on and off fairly rapidly it's not going to happen.

kc: I don't see how to make that work in the context of the current system. I think it's an intriguing concept, and I applaud the vision. It's just that we don't have a framework to put that into right now. And, he hasn't proposed a way for us to get there that I've heard of. But, again, I would be very happy to be wrong about all this.

Let me sum up what we have done with our Internet core mapping project. The poster shown at <http://www.caida.org/analysis/topology/as_core_network/> describes the methodology and draws a picture. It also has an interactive java applet that lets you navigate around this core of the Internet. It labels all of the autonomous systems, describes what it means to be core, etc. Other pieces of the topology mapping project that we're looking at are ways to identify critical pieces of the infrastructure. Not necessarily core of the Internet, but critical pieces of a particular intranet. For example an ISP might want to know which pieces of their infrastructure have the least amount of redundancy. We're interested in also doing long-term trends workload analysis. I'm sent several questions a week asking me long-term trend questions about workload, which I have a hard time answering because we don't have very much long-term data. We do have a paper we submitted for publication in September, which I can't give you the URL until it is published, but it describes some long-term trend analysis over the last year at the FIX-West exchange point.

COOK Report: How do you define long-term?

kc: 'One year' for this paper because that's all the data we had. Another big problem right now for both research and useful tool development is finding visual metaphors for presentation of Internet data. When you end up with over 30-60,000 trace routes and you want to depict macroscopic insights into the infrastructure, there aren't really good basic visual models to do that. Consequently we're having to come up with our own. Here I'm talking about a set of techniques, not a given tool. I'm talking about ways to present data in general, so coming up with new ways to depict data. I don't know if you've read

much of Edward Tuft. Scientific Information Visualization Techniques. It's not a given tool. It's just a way to depict data. We've had to come up with several different ones in the course of our investigations because there's too much data to explain by simply drawing a map.

Moore: You don't always know what you're looking for when you start, and visualization techniques let you see things sometimes that are present in the data, patterns in the data, that aren't obvious the other ways of looking at it. So, we develop visualization techniques both to ask specific questions, and to look at data in new ways that enable insights you wouldn't have found otherwise.

kc: There really isn't an enormous amount of data in the Internet right now. Different people are doing different kinds of active measurement, passive measurement, putting traces up so that other people can do research on them. There's really little concerted efforts in correlating different data, both similar types of data from different places, similar types of data from the same place over different times, and different types of data. Correlating, for example, active measurements of latency and round trip time with routing updates or workload changes, or topology changes.

COOK Report: The better correlation will lead to the ability to understand how to recover from a fiber cut or a route flap?

kc: Well, more generally, if you ascertain a correlation between two kinds of data, and you have one, you can make inferences about the other.

Monk: This is very important as to where Caimis will be going. One of the critical components of network management systems is collecting lots of different types of data and correlating some of the data together, and you have the connections linked within the network management system.

COOK Report: I'm assuming that's one of the things related to the multi-gigabyte and terabyte routers that are waiting in the wings. If I understand it correctly from talking to Ironbridge last summer, you can't really use a terabyte router in a meaningful way without a very good network management system.

Monk: There are lots of things that are coming on the market now and technology is becoming increasingly sophisticated to the degree that the need for monitoring and management is increasing on all levels.

Part Two: Network Management as an Example of Caimis' Mission

Monk: We should start by stating that Caimis' mission is to be the premier provider of monitoring, network management, and geographic localization solutions for optimizing planning, engineering, and transactions at the provider, enterprise, and user levels. Caimis will focus on three areas. The first is monitoring, and we have a variety of solutions that are for passive and active monitoring through active polling and passive analysis of router and link data. The second is network management systems for acquiring, storing, correlating, and analyzing data. The third area, geographic localization services, involves a set of solutions aimed at mapping different types of network data to geographic locations. It is the latter area that has direct relevance for e-commerce and financial transactions, as well as providing a different type of visibility into the infrastructure that can potentially be very relevant to providers and users. [Editor's Note: In June Caimis formed an affiliate company, Caimis Geo, Inc., to develop the geographic location services business. This will permit Caimis, Inc. to focus on monitoring and network management systems.]

COOK Report: You spoke of "transactions." A transaction is what is happening between two users or two systems, and where it is happening as well?

Monk: A couple of different things here. One is being able to identify the location of end users, so if you are browsing the web, a content provider would have a better idea of where you're physically located. Such knowledge would enable it to provide you with content best suited to your geographic location. For financial transaction purposes, having the ability to identify the location, whether it is at the country level, state level, city level, of the person conducting the transaction is very powerful verification information.

COOK Report: And, one reason it's powerful is that it is useful in capacity planning, for example?

Monk: That's a different purpose. That's what an ISP would use this information for, not what companies engaging in financial transactions and supporting financial transactions would use it for.

COOK Report: Okay, and their purpose would be what? Just understanding where their customers are?

Moore: It's primarily aimed more towards content providers or e-commerce websites, where they want to know where their customers - perhaps because they have restrictions on distribution of software to people in different region, geographically based or by country. Or, websites that might want to change the language presented on their homepage depending on which country the user is coming from or return targeted content based on where the user lives, e.g., content appropriate to Southern California versus Boston.

COOK Report: So much for CaimisGeo, Inc. then. Please explain why Caimis' two areas — monitoring and network management systems are of growing importance.

McRobb: Caimis believes that network management is becoming increasingly important in the Internet. One of the main reasons for that is the push for business-to-business communications over the Internet. All the research firms are predicting that the business-to-business market is going to grow by a very large amount in the next three years. That could turn out to be false, because it's just research or predictions. I tend to believe them though. Let's take as an example, the Automotive Network Exchange (ANX), large industry consortium of the automotive corporations in the US found that two years ago when they went out looking for a solution based on the Internet for their business-to-business communications, they couldn't get it. None of the providers even within their own network could provide the level of service required by the ANX.

About four years ago, Bob Moskowitz, the representative from the ANX at the time, came and gave a presentation at one of CAIDA's workshops on the requirements of the ANX and what they were looking to do. He was practically booed off the stage by providers. He was told by most of the engineers in the audience that were representing various large providers that basically they weren't interested in his business and there was no way they could meet his technical requirements. This was unfortunate given the potential money involved in the business-to-business market.

COOK Report: I did a couple of interviews, but it's been about two years, and I'm not aware of how well the Automotive Network Exchange has worked in practice once they did get up and running with a few certified providers. Are you aware of that, or is that relevant to what you're saying or not?

Daniel: I don't think that's relevant. I'm using the ANX as an example of the business-to-business market and the difficulties that they encountered are the same difficulties that other people that want to do business-to-business communications encounter.

COOK Report: If ANX had found a wonderful solution, we would have heard about it by now, presumably?

McRobb: Yes. There are several issues here. One is that on the business side of things it is difficult to do these things inter-provider today. There's lots of legal entanglements, there's really no industry agreed parameters to be passing between providers for things like SLA guarantees. Caimis' notion is that network management is becoming increasingly important because the customers are demanding a higher level of service. Today a lot of the Internet business remains business-to-consumer. There is still a considerable chunk of business-to-business networks that are not using the Internet. This will probably continue to be the case until there are solutions that allow providers to meet the service level requirements of the business-to-business community.

Interprovider Service Level Agreements

McRobb: So, today most providers actually have fairly weak management systems in place within their own network. Unfortunately, network management and IP space has sort of been in a sad state for as long as I've been in the business, which goes back to 1991. Until that situation improves, I think the business-to-business market it going to be hard to realize. Caimis' hope is to provide some solutions that will help providers address the requirements of the business-to-business market.

COOK Report: The reason that this is still difficult is that to get the total geographic flexibility where, if I'm a business, I can get to the location anyone else on the face of the earth, I'm going to have to cross a network provider boundary. If I'm a UUNET, I can say to anyone as long as you're a UUNET customer I can promise you wonderful service, but if I'm a UUNET, I cannot promise GTE or Sprint or other people the service I can if they're on my own network, right?

McRobb: It's actually a combination of both. Today, if you look at what's in the provider's service level agreements, it's many, many miles away from what a customer expects.

COOK Report: Even for service within its own network, never mind crossing a network boundary?

McRobb: Yes. There are a variety of things at play there. One is that network management technology is still fairly nascent in the IT world. Two is that most of the companies providing network management solutions don't really understand the problem domain.

The people like Hewlett Packard who are producing network management products tend to be the largest NMS vendors. HP, Cabletron, Seagate, some of the other ones, they don't really understand the problem domain. They have these vertical applications that are fine for element management. They don't connect into overall network management. They also have almost no interfaces with service areas associated with running major networks.

COOK Report: When you say element management, did you mean something like local area network management?

McRobb: I mean something like the devices. Currently solutions can tell you what's going on with a particular router, but when you need to move up to something like defining an end-to-end service for a particular customer, there is a big void. Providing such solutions are challenges that happen both within and between providers.

When you start talking about inter-providers, the situation gets worse. Information can't be exchanged. Providers don't want to provide any sort of service level guarantees for infrastructure that they don't control and to which they have no access. This is sort of the state of things today. Our hope is to provide monitoring solutions in that space, in addition to raising the bar for network management in IP space.

What I mean by raising the bar is a variety of things. One is to come up with some sort of common interfaces that providers can use within their own network, but between different systems. When I talked about the element management versus service areas, that's a big problem today.

COOK Report: You said between different systems. A different system would be what?

McRobb: For example, provisioning and network management itself, element management. Today, basically most of the network management systems sit in the element level, meaning they tell you what's going on with particular devices in the network. They don't interface well with any of the other business processes within a provider.

COOK Report: This is processes within things such as the traffic measurement, billing systems?

McRobb: Yes, customer care processes, billing, provisioning, ordering, trouble-ticketing, which is a big problem in providing service to the business-to-business environment. A good comparison is if you look at something like the telephone service. Today you can call up your local RBOC and order a whole variety of services internally within the telephone company all of their systems

communicate with one another. It's feasible for an RBOC, for example, to provide a webpage where you go in and fill out a form, order all the services you want, and everything else is kicked off automatically. Including customer care processes, the actual provisioning, the monitoring, keeping track of whether or not the service is actually meeting the service level you're expecting, etc. In Internet space those things rarely happen, which is unfortunate. There's a big need for that. We're hoping to address some of those things. We're not going to tackle the whole ball of wax. We don't have the resources to do that.

COOK Report: But, if you're successful, hopefully you'll build them up along the way.

McRobb: Right. One of the key things for us to do, as we're looking at these technologies that we license from the university and as we develop our own technologies, is to keep an eye glued to the big picture.

COOK Report: The big picture is general protocol development, general commercial needs, general operational conditions and problems?

McRobb: The big picture is all of the business processes involved in managing a network, which would include things like provisioning, customer care processes like trouble ticketing, fault reporting, and performance monitoring. Today most of the network management companies don't address any of the business processes other than element management in Internet space. If you go over to the telco side, it's a different story.

COOK Report: One of the most critical areas is cross provider issues, quality of service issues, and differentiated services. That seems to be a holy grail that a lot of people are thirsting after, but the progress isn't too huge. You eluded to large providers being reluctant in these areas because they consider their processes or databases proprietary. They don't want to share this, that or the other thing with their competitors.

Has it occurred to any of you all to try to find out whether Caimis could become a trusted neutral third party or intermediary. A service provider could take its proprietary data to Caimis and a competitor could also plug in its data. You could do the necessary linking and translating without giving a way from one customer to another customer the secrets the customer want protected?

Monk: No. Caimis will do consulting relating to customized solutions for large providers and large institutions. We will then develop standardized products that can then be sold throughout the industry — at the provider level, at the small ISP level, and to

institution users. Caimis will not provide services at the level that you're discussing. That is not part of our business model.

Lessons at the Interface Specification and Software Implementation Level

McRobb: Our involvement in terms of things happening between providers is only at the interface specification and software implementation level. We will not be providing services. We will not be storing provider's data for them. We will not serve as an exchange point between providers.

COOK Report: Can you give me an example of what you mean by interface specification? In other words, a provider might say I'm looking to get my data out in certain ways and I want it in gigabytes rather than terabytes?

McRobb: Yes. That's basically the case. When I say standardized interfaces, I mean within a provider where they have sort of a common set of methodologies and API (meaning application programming interfaces) to our systems and other systems, and, also intra-provider. You will find a good case of this if you look at what's happened at the telecommunications industry but not in the IP industry. The whole deregulation of the industry forced the telecommunications providers to do a considerable amount of work on that front. There's a lot of lessons to be taken from there, and we will be applying some of those lessons to the Internet.

COOK Report: It's called Telcordia now, but before it was called Bellcore. Ten years ago or longer, did Bellcore do for the phone industry some of the kinds of things that you are talking about Caimis doing?

McRobb: Bellcore really didn't, the ITU did.

COOK Report: You're talking about the International Telecommunications Union ?

McRobb: What I'm saying here is that the actual deregulation of the industry is what forced a lot of these things to happen, and before that most of the solutions were actually proprietary because there was one giant telephone company.

Moore: The idea here is to define a standard way so that people can exchange information basically between different ISPs, that it's in the same format, the information means the same thing to both people, they're talking the same things, and the software can talk to each other, too.

McRobb: There are a variety of industry consortiums that have looked at this problem, looked at it from a bird's eye view. There are a few that have looked at things a little further down.

COOK Report: And, when you say industry consortium, you mean Internet?

McRobb: No, which is unfortunate.

COOK Report: IOPs, which was under the XIWT at CNRI was a potential group that could have been in this area. Have they done much of anything?

Monk: I think you'd have to talk to Chuck Brownstein at XIWT to get a clearer picture on that. My understanding is that they've done a little bit, and I don't know how much.

COOK Report: If problems had been totally solved, you wouldn't be talking about working on them.

Monk: I believe that IOPs is working mostly within the confines of the ISPs that are members of XIWT.

McRobb: One of the things you have to consider is that IOPs is not an open forum.

COOK Report: Are you talking about playing a kind of IETF role in helping develop some open standard for defining interfaces?

McRobb: Yes and no. We don't intend to become an IETF. But, yes in terms of working within and working with various industry consortia. Yes, we will work with the IETF, OMG, JIDM, NMF, and those are the sorts of consortiums that have already done some work in this area.

COOK Report: JIDM?

McRobb: JIDM stands for Joint Interdomain Management (<http://www.jidm.org/>). The NMF (<http://www.nmf.org/>) was the Network Management Forum, it's now called the Telemangement Forum. The OMG is the Object Management Group (<http://www.omg.org/>), they are the ones that defined CORBA, middleware technology. They have a telecommunications operations group. They have done some high level, and actually some low level, specification work, applicable to telecommunications network management. There's another group called TINAC (<http://www.tinac.com/>) which has been working on defining specifications for network management in the telecommunications domain. They haven't done a lot of work in IP space. In fact, most of these groups have not done a lot of work in IP space, but there are an awful lot of valuable lessons to learn from the work that they have done, then apply it to the Internet.

COOK Report: You're doing a very good job from what I would regard as a middle to longish term point of view as far as defining a set of needs and problems and defining Caimis as providing some solutions to the problems. Once you get ready to be announced and take on business, can you help me understand what this means to both my commercial readers and customers? I have more and more customers now in the financial, investment bank, stockbroker areas that are trying to define what kind of impact things are going to have on the net. How does one understand in the first six months of Caimis' existence what Caimis offers in the short term, or is it really more of a middle to a long term solution?

McRobb: No, we do have solutions that we plan to make available in the short term. For example, we have some technologies for things like reachability and performance monitoring. We have some ICMP-based tools that do things like monitor round trip time, monitor packet loss between a source and any number of destinations. For example, we are currently running some of our tools against approximately 75,000 destinations.

COOK Report: If I understand what you're saying, there are tools now that have been developed under CAIDA that are in existence that ISPs know about, but they are reluctant to really start using seriously because then they would mean that they are depending on them. What you're saying is that from the word go, for the ones who want to start dealing with this, we're going to talk about the following kind of tools, or tool sets, or programs. We are here and available now to serve as your source of support for use of these tools for commercial purposes.

McRobb: Correct. We're not going to tackle the whole picture. We have particular technologies that we can apply today for particular areas. Our hope is to always keep those tools in a state where they are amenable to being integrated into the big picture.

COOK Report: Assume I'm John Curran with NexLink responsible for operations network, what are the tools that Caimis would have to offer John, and theoretically at least, how would they solve some of the problems that a new greenfield provider like a NexLink might have? Is that a reasonable question?

McRobb: Yes. In a general case, tools that we have for large providers would be things like skping and skpingd, which are reachability tools.

COOK Report: Who is involved with Caimis and do you have financing?

Monk: Our board of directors includes kc, Daniel [McRobb], myself and David Rand who is co-founder of AboveNet Communications and senior vice president and Chief Technical Officer at Metromedia Fiber Network. We are currently operating on seed funds provided by board members, but anticipate a round of financing in August.

Caimis: Operational and Mission Structure

COOK Report: I have the impression that you see Caimis as being a vehicle that you can use as a means of providing desirable professional employment for some very, very good and knowledgeable (I think you mentioned some old ANS programmers) who otherwise might be tempted to go on in different directions. Are there any of those kind of staffing issues?

Monk: I didn't touch on this and I think it's important to note that Caimis will have two offices. The San Diego headquarters will focus on the development of the geographic-locations services under Caimis Geo, Inc that I mentioned earlier. The San Diego headquarters will also be doing all of the productization, sales and distribution of all of our monitors and network management systems. The Ann Arbor office will be the leading edge site for senior software and network engineers who are doing customized network management solutions. They will be the ones staying on the cutting edge technology-wise and requirement-wise for the large providers. What we will do in San Diego is take some of these solutions and develop standardized products that can then be distributed and made available for the smaller ISPs, enterprises, and others.

COOK Report: So Ann Arbor gets with the big guys and some of the open source packages and develops those into the kinds of things that the San Diego office can productize for some of the smaller providers. Am I right?

Monk: And, also large institutional users. Because the enterprises and large institutions are as much of a market for this as the small ISPs.

COOK Report: For example, if I'm Chrysler or Ford or GM, and I have manufacturing plants around the country, I've got networks of my own for which these are useful?

Monk: Absolutely, as well as nearly all your Fortune 500 companies which are interested in having more insight and visibility as to the performance of their WANs as well as the performance of services they are receiving from their ISPs.

COOK Report: From the point of view of a Fortune 500 company, would these tools be as applicable to what they would want to do on the public Internet, would they also be applicable to measuring and working with their performance characteristics of their virtual private networks?

McRobb: Both. We have tools that would be effective in both cases. In fact, probably for some of these tools, the VPN market is one of the customer bases that will be very significant from the beginning.

COOK Report: Sounds good. Any reason that somebody hasn't done these tools for the VPN market before?

McRobb: There are some companies that have done some tools in this base, including some of the NMS vendors, like Hewlett Packard. They have a product called Firehunter.

COOK Report: But what you're offering here are tools that are going to be more open sourced, people with knowledge both of the Internet and intranets, and something that presumably can be more attractive because it's more leading edge and more flexible and covers all the worlds.

McRobb: It won't be open source. That's sort of a loaded term. Everyone has their different definition. When I think of open source, I think of GNU public license, but we will provide source code to large strategic customers. I think the significant benefit of Caimis is that all of the employees that are currently slated to work in the Ann Arbor office have significant network management experience, meaning they've actually managed large networks. Most of the other software companies doing this work don't have programmers who understand the problem domain — the Internet. As I mentioned earlier, one of the big problems in this space today, is that the network management vendors, even the ones who are particularly good at things like element management, don't understand any of the other business processes within a provider. However, they really have to integrate everything in order to provide an effective solution.

COOK Report: So is there anything else that you can tell my readers about Caimis' plans?

Monk: We'll announce Caimis in July and plan to start by offering a simple UNIX desktop monitoring tool named skping that is used for visually depicting latency and loss measurements as well as displaying the results of frequency analyses done on the round-trip time measurements. Provider and enterprise products will be available in September. That's probably enough for now.

Executive Summary

CAIDA - Caimis pp. 1 - 8,

We interview Tracie Monk, Kim Claffy (kc), David Moore, and Daniel McRobb on the details of CAIDA's creation of Caimis. Caimis www.caimis.com is a profit making spin off from CAIDA that will take many of CAIDA's research oriented, network traffic management tools and upgrade them to production quality commercial software for use by major ISPs. Kc explains the CAIDA Internet Teaching Laboratory program, the San Diego Nap and plans for OC48 mon - as well as CAIDA's work on defining and understanding the Internet core.

The Caimis web site went live on July 12. This interview is the first detailed description of Caimis to be published. Caimis will focus on monitoring systems, network management systems and geographic location systems. In the monitoring area it will offer a variety of solutions that are for passive and active monitoring through active polling and passive analysis of router and link data. Its network management systems will be used for acquiring, storing, correlating, and analyzing data. The third area, geographic localization services, involves a set of solutions aimed at mapping different types of network data to geographic locations. It is the latter area that has direct relevance for e-commerce and financial transactions.

In June Caimis formed an affiliate company, Caimis Geo, Inc., to develop the geographic location services business. These will be tools aimed more towards content providers or e-commerce websites, where they want to know where their customers are actually located - perhaps because they have restrictions on distribution of software to people in different regions, restrictions that are geographically based or by country. Or, you might have web sites that might want to change the language presented on their homepage depending on which country the user is coming from or return targeted content based on where the user lives, e.g., content appropriate to Southern California versus Boston.

Caimis finds that ISP service level agreements generally still do not measure up to the standards that business to business customers expect to see. Hewlett Packard, Cabletron and Seagate produce network management systems but do so without an adequate understanding of the Internet problem domain. While current solutions can tell network engineers what's going on with a

Continued on page 25

Cogent Communications Guarantees 100 Mbs Internet for \$1000 per Month Buys Dark Fiber and Builds National Network Optimized Entirely for Data - Business Model Emphasizes 100 BaseT LAN Interconnection in Buildings it Serves

Editor's Note: Dave **Schaeffer** is CEO of Cogent Communications, Inc. He grew up in the Washington DC area and earned a BA in Physics and a PhD in Economics. He has developed numerous communications enterprises including an SMR carrier and paging carrier. His most recent venture prior to Cogent was Pathnet which was an aggregator of excess telecommunications capacity from multiple utilities around the country. We interviewed him on June 16, June 19 and June 23.

COOK Report: In a sentence or two how did Pathnet work?

Schaeffer: Pathnet used a combination of fiber optic facilities and SONET microwaves. We upgraded various disjoint utility systems and then integrated them into a national network. We would get part of their fiber for having done this for them. We then used the fiber to sell wholesale telecommunications capacity to any provider whether it be voice or data. ATT and MCI WorldCom were our largest customers. We were spread out over the entire country.

COOK Report: How then did you get into the launch of Cogent?

Schaeffer: We realized both that there was tremendous amount of growth in Internet access demand and that there was new technology that would drive the cost per bit mile of traffic down significantly. It is also our belief that demand is relatively elastic and that if we drove prices down, customers would find additional use for the bandwidth. The idea was to capture a set of technology changes and market changes that allowed us to build, for the first time, a facilities-based infrastructure solely for the purpose of driving Internet traffic. This will be in contrast to the way in which traditional ISPs have built their networks which was to deploy routers and then rent capacity between those routers on a network that was designed primarily for voice traffic with a data overlay.

COOK Report: In other words you are going to be one of the first to take advantage of the kinds of technology making TCP/IP over fiber available at extremely low cost that was used in the adoption of CANet3 two years ago. The key seems to be expansion

of bandwidth. I saw an assertion recently that for every one per cent decrease in the price of bandwidth demand grows by 1.5 percent. What is your point of view on this issue?

Schaeffer: Bandwidth growth in the Internet is in unit terms somewhere between 200 and 400 per cent per year. Existing telco-voice oriented infrastructure is simply not going to support this kind of growth cost-effectively. Even the SONET infrastructure that is built for the joint purpose of voice and data is very expensive when compared to our network that is optimized purely for data solutions. With this infrastructure we are able to drive down the cost of data delivered per-bit-mile by two orders of magnitude over where existing networks operate.

COOK Report: How did you come to realize that this opportunity existed?

Schaeffer: It was a combination of several factors. In particular I listened to a speech by Michael Armstrong in which he said that ATT had the lowest data cost in the industry even though it had an antiquated network with a high cost structure. ATT was able to do this with its branding and marketing was able to drive its network to a greater percentage of utilization that those of its competitors. Because networks are a fixed-cost asset, the hidden costs were substantially lower.

So as I looked at the current state of infrastructure and concluded that I should build a facilities-based Internet backbone where there was tremendous pent up demand for high bandwidth application. Technology was making it possible to drive significantly higher utilization from existing infrastructure — no matter what the state of the earlier network. It was a combination of factors ranging from the ability to purchase dark fiber, to the ability to go after traffic that was the fastest growing - namely the data traffic.

Now on the revenue side, the reasons why networks produce a lower revenue-per-bit-mile for data than for voice is that a bit of voice travels a mile at *seven times* the revenue of a bit of data on the American network. In fact American networks today derive 85% of their total revenues from voice,

yet less than 50% of the volume of their traffic is generated by voice.

Acquisition of Infrastructure

COOK Report: From an article published yesterday (June 15) in Light Reading (http://www.lightreading.com/document.asp?doc_id=778) you have some impressive figures on infrastructure purchases. "Local dark fiber rings from Metromedia Fiber Network: \$100 million. Nationwide long-haul fiber from Williams Communications: \$215 million." What exactly did you buy from Williams?

Schaeffer: We bought a 30 year IRU on an entire strand of Williams fiber which gives us complete control over the fiber for its useful lifetime. As DWDM continues to improve, we can, at will, over the life of the agreement change the optronics any time we desire to. We actually purchased two strands forming a single pair throughout their entire network. Our strands do have physical diversity but, as far as a cut is concerned, our network is just as well protected as theirs since any cut will take out the entire bundle and not just one or two strands. We always have two paths two every site. We also pay a fee to Williams over the life of the agreement to maintain, surveil and repair the fiber should there be any cuts. There are significant penalties should a repair by Williams fall outside the parameters of our service level agreement.

COOK Report: How are you dealing with the issues and expense of SONET protection?

Schaeffer: We are protecting our network at Layer 3. Therefore a consistent way to think of the Cogent network is as a facilities-based nationwide switched LAN.

COOK Report: How is what you are doing different from someone who is running TCP/IP over glass?

Schaeffer: This is what we are doing but we are also eliminating the SONET protection layer from our network. We are not with holding an entire separate strand of fiber

from our network for SONET restoration in the event of a fiber cut.

COOK Report: By forgoing the SONET you double your amount of available fiber. Don't you need SONET to do your optical to electrical conversion?

Schaeffer: We are not buying any SONET equipment. None what-so-ever. We are doing optical to electrical conversion but we are doing it predominantly within the terabit class routers. We are running packet over SONET framing. Gigabit Ethernet is used in the local loop. SONET framing in long haul. The SONET framing is a transport standard that is built into the Cisco routers. Consequently there is no additional cost. Since we are doing our protection at Layer 3 using the router, we do not need the redundant protection from SONET. Instead of 50 milli-second recovery, we get one second recovery at the router level which, for data, is adequate. Any physical disruption in service would cause the core routers to start rerouting packets. This process takes about one second.

Because we do not oversubscribe, we are at a maximum two to one oversubscribed if we take a fiber cut which is a lot better than a traditional ISP which may be as much as forty to one over subscribed.

COOK Report: How do you figure your provisioning? In other words for every 100 megabit-per-second customer how much bandwidth do you provision?

Schaeffer: Every one who buys a hundred, gets a hundred. We don't cut corners. Every one gets a full hundred-megabit connection all the way through the network.

COOK Report: Then if you have an OC192 between New York and San Francisco, how many customers would you place on such a link?

Schaeffer: Since an OC-192 is 9.6 gigabits, it would support 96 customers. Because the packet will get on and off the network multiple times, we assume that the average packet will travel 1500 miles on our backbone.

COOK Report: Then you do some degree of multiplexing?

Schaeffer: No I disagree with your judgement. We observe actual traffic and then make predictive assumptions based on our observations.

COOK Report: But if you had ten customers for a one gigabit, three thousand mile haul and, if the average packet distance were 1500 miles, you might feel you were giving them the full bandwidth if you put twenty

customers there? This is what I meant by multiplexing.

Schaeffer: OK. Yes that is what we would do. I would say that our network offers inherently the same reliability as any other network. However, because we offer much greater capacity, we can assure that a much higher service level agreement can be established through the fact that we do not oversubscribe and that we therefore do not drop packets across our network.

Network Hardware

COOK Report: Can we try to take a logical progression of discussion to understand your network build out starting with your acquisition of your fiber? How do you choose a fiber provider both in metropolitan areas and across country?

Schaeffer: Our selection of fiber vendors was based on several criteria. The first being the geographic reach of their networks, both on the long haul inter-city networks, and the local intra-city networks. Secondly we examined actual deployment versus proposed deployment. We assessed each vendor in terms of how much of their network was actually completed.

COOK Report: What do you expect your geographic coverage to be? How many cities?

Schaeffer: It is our intention over time to widen our geographic coverage based on fiber availability and on our successful execution in the major markets in which we begin operation.

COOK Report: A very extensive network would certainly be highly desirable then. Williams has one of the most extensive. How did you settle on them?

Schaeffer: We assessed multiple providers on the basis of completeness of their networks, and the actual quality of their glass. We looked at shelter spacings and at the condition of their shelters.

COOK Report: What do you mean by shelter spacings?

Schaeffer: For the long haul network one must look at the spacing between each of the transmission sites which would house either an electrical regenerator or an optical amplifier. It is a specialized form of co-location that is totally defined by the physical route in which the provider from whom you are buying fiber has these shelters. We really have no alternative or siting gear in the middle of a cornfield in Iowa, other than going to whomever is selling us the fiber that we have chosen to buy.

COOK Report: So the shelter is where you would need to install your transport amplification equipment?

Schaeffer: Right. We would expect our network provider to supply those shelter co-location facilities to us. However different providers have different distances between their network shelter spacing.

COOK Report: How do you decide what is good when you are talking spacing?

Schaeffer: Some of the new technologies require more shelters so that you can place amplifiers closer together and minimize the number of electrical regenerations.

COOK Report: Does your interest in avoiding SONET (except for SONET framing) have an impact on spacing decisions?

Schaeffer: It really does not. A spacing decision is determined, at least in part, by the gain in amplifiers and the power level in those amplifiers as well as the launch power in the transponders. In evaluating different vendors, one must look not, just at a simple one dimensional metric but rather at all of these different criteria in trying to establish the optimal network.

COOK Report: So you are saying that you better know all the equipment configurations in order to figure out what shelter spacings will be appropriate?

Schaeffer: Yes. There are some vendors whose equipment will only work with certain physical shelter spacing. As we evaluated our potential dark fiber vendor, we looked for that network vendor whose shelter spacing would support the maximum number of equipment vendors. Note that the shelter spacings are not inter connect points. They are there only to amplify the signal. All that we need to be able to do is to rent standard size rack space in the shelters.

COOK Report: What do these shelters look like?

Schaeffer: They would typically be about ten feet tall, twelve wide and thirty feet long. They would look like a trailer.

COOK Report: And you could put these anywhere that you really needed?

Schaeffer: That's correct, provided you can get the appropriate permitting and that you could bring in electricity and a concrete foundation. Williams ranked very high in all these categories. We used a multi-dimensional matrix to evaluate each of our vendors. It is fair to state that we did an exhaustive evaluation of all providers. Both existing as well as proposed and both on the long haul as well as for the intra-city.

COOK Report: In addition to the spacing and geographic coverage, what did your matrix cover?

Schaeffer: Willingness to Sell and the Pricing.

COOK Report: Besides Metro Media who offers competition at the local level?

Schaeffer: Level 3 is a local seller. As well as there usually exists in each market one or more local providers. They could be utilities, municipalities, state governments, cable operators, or a CLEC. Ultimately, at the local level, you also have the fall back position of going to the RBOC. Consequently there is a wide variety of sources of potential fiber.

COOK Report: If you want to be able to deal with only one company in most areas, does this limit you to Level 3 and Metro Media?

Schaeffer: Yes, but there is also a new local fiber start up out of Chicago called Looking Glass. While it does simplify things to choose a national player that is in all the cities to which we want to go to, we have found that no single local vendor goes everywhere that we would like to go.

COOK Report: At the local level then what are your selection criteria?

Schaeffer: You would look at the physical topology of the rings, which buildings are going to be served, the type of fiber that has been deployed, the span budgets and splice budgets or that fiber.

COOK Report: What are the span and splice budgets all about?

Schaeffer: They determine how far you need to go before you can optically re-amplify or regenerate your signal. What is involved is the type of fiber - its overall quality as well as its insulation. All dark fiber is by no means the same. Fiber can be optimized for the number of wave lengths, for distance or for the power applied to it. There are a number of different design criteria and each vendor takes a slightly different approach to which of those criteria they emphasize with their fiber.

COOK Report: So for example, given one set of optronics you could use fiber from different vendors and find that you could get differing numbers of wave lengths of light?

Schaeffer: That is correct.

COOK Report: How do you interconnect your wide area and your local fiber?

Schaeffer: Typically there are a number of

interconnect points from both of the networks, and yes there is an engineering challenge, but we are designing our networks so that the local and long haul directly interconnect with each other.

Selection of Optronics and Routing Equipment

COOK Report: How did you put all this together? When did you begin your planning?

Schaeffer: In the fall of 1999 we spent almost four months dividing the network that we are building into four areas of technology. We then methodically reviewed all of the equipment available in each of those areas. They were the long haul WDM equipment, the terabit routers, the metro WDM equipment, and the edge gigabit routers. We evaluated the four technologies by the use of four criteria: the technology; the availability; the price; and the ability of the vendor to support the product.

COOK Report: How did you prioritize these criteria and how do they affect each other?

Schaeffer: We choose technologies that were all based on industry standards in terms of interfaces and interconnectability. We were not necessarily married to a single vendor. We felt very comfortable going with the four different vendors for the four different technologies. Each one was evaluated independently of the other.

COOK Report: Could we start with the long haul WDM equipment and the routers for your long haul backbone?

Schaeffer: For the wave length technology we looked at Lucent, Siemens, Nortel, Cisco, Ciena, Alcatel, Qtera, and Sycamore. We selected the Cisco. (Actually Pirelli, a company just purchased by Cisco.)

For the terabit routers we looked at Nexabit, Avici, Pluris, Cisco, Juniper, the Netcore product that is now owned by Tellabs and at the Argon box. We applied our four dimensional matrix to the products. Many of the products were not yet real. There was more vaporware than real boxes that actually are available and had all the features and functions we need — including stabile protocols and routing algorithms.

COOK Report: Certainly Juniper is very real competition for Cisco. What happened that knocked Juniper out of competition with Cisco?

Schaeffer: The Juniper switching fabric was smaller than that of Cisco. Eighty gigabits to 320. Both had stabile BGP 4 implementations. We felt that Cisco's IOS manage-

ment software was more stabile than Junos. Both vendors had the same inputs and outputs. The ability to fully switch all of this on a non blocking wireline speed basis was greater in the Cisco product than it was in the Juniper.

COOK Report: What are you going to be switching and why is having a lot of switchability good for you?

Schaeffer: We need to make sure that, at wireline speed, for a reasonable distribution of packet sizes 100% of the packets would go through the switch even when configured for the maximum number of in and out speeds. We want this full switching power so that we do not ever run the risk of doping packets. Most networks try to force more packets through than the switch can ultimately handle and, as a result, the switch discards some of those packets.

COOK Report: Then what you are saying is that with the switching you want the greatest possible number of paths so that like a car on a freeway at toll booths your packets never have to risk slowing down in order to get through.

Schaeffer: Yes.

COOK Report: And the Metropolitan WDM equipment?

Schaeffer: We looked product by nearly 15 different vendors. Companies such as Nortel, Lucent, Ciena, Chromatis - which we chose for a small portion of our network. We looked at Alidian, at Cyrus, at Astralpoint, at Kestrel, Ignitus, Serrento, and Cerent (which is now part of Cisco). We evaluated these vendors on their technology, and availability. In particular we looked at the number of wave lengths each system could support, its throughput, its price-per-bit and the way in which it handled the optical add drop. There we chose a predominantly Cerent solution.

COOK Report: According to Light Reading: "Chromatis offered a way for Cogent to split its leased fiber into DWDM channels, which then could be aggregated and switched as needed throughout the metro network." Why did you wind up with both Chromatis and Cerent?

Schaeffer: Chromatis has grid spacing that is more applicable to dense fiber rings than Cerent. In other words the Chromatis product can support more wave lengths than the Cisco products. With some metro rings the lower wave length count is acceptable. Other metro rings need greater capability.

COOK Report: I assume the rings we are talking about belong to Metromedia Fiber which has different number of rings in dif-

ferent cities. Do they say just what they do offer in any given city they are in?

Schaeffer: No. They keep that information proprietary and confidential.

COOK Report: What might be a reason for the number of rings in cities of similar population to vary from each other?

Schaeffer: I'd guess it would depend on the building density and the city's physical topology over which its buildings are distributed.

COOK Report: And what do you have to give up for those extra wavelengths from Chromatis? Anything much?

Schaeffer: Just more dollars. Availability, price and functionality were key components. Chromatis also had some very flexible features in terms of configuration and use that were somewhat more unique than some of the more limited products — for example the add drop capability, the number of cards that could be supported in its chassis, the actual backplane switching fabric and the speed of the cards. The range of capabilities that can be applied over a range of installations is extensive.

Now finally for the edge routers we looked extensively at the Acelor product from Nortel, at Cisco, at Foundry and Extreme and we chose Cisco's 2948G-L3.

COOK Report: And what was guiding your decision making?

Schaeffer: The performance of the packet forwarding engine, the feeds and speeds, the price per port, the protocols supported, the management systems, the physical size, the power requirements.

COOK Report: And management systems would be what?

Schaeffer: IOS. The rough equivalent to a PC operating system.

COOK Report: Some people advocate, (Ross Callon at Ironbridge for example) that a network management system for a high speed router be terribly mathematically sophisticated at doing things like looking at traffic flows to permit traffic engineering and that given the amount of traffic to watch and, given the time in which to make decisions, they suggest that there must be a lot of artificial intelligence involved. In looking at your own network management did you find out that you needed to face these kinds of issues or did you find that you could get around them by throwing bits at the solution as it were?

Schaeffer: I don't think it was so much the

complexity as the ease of use that we looked at. We did carefully examine how stable the management systems were and how long they had been in place. It was a holistic view that looked at all four criteria and again we choose the Cisco solution but we were perfectly prepared - as a matter of fact we expected to have to buy technologies from multiple vendors. In fact we were very surprised that in our evaluative process, Cisco won in multiple categories.

COOK Report: In the *Light Reading* article on you they wrote: "Cisco had other things going for it. It was able commit to working with Cogent on consistent upgrades, and to offer reasonable financing terms. And Cisco could promise Cogent a guaranteed level of marketing in exchange for becoming a so-called "Cisco Powered Network"—running a specified (and undisclosed) percentage of its network on Cisco gear alone. Cogent signed a three-year, \$280 million deal with Cisco on March 6." What is your reaction to this assessment?

Schaeffer: I think we did not really need this special designation. It really did not enter into our evaluative process.

COOK Report: But help me to understand what it means to be a Cisco Powered Network.

Schaeffer: This means that we have certain content requirements that our network needs to meet in terms of Cisco equipment. It also means that Cisco provides additional engineering and design support so that we meet certain quality levels so that when we sell our service to our end user customers they can be assured that we meet end user standards.

COOK Report: Again according to *Light Reading*: "Each GSR [editor: Cisco edge gigabit router] generates multiple gigabit Ethernet links that feed into multiple Cisco ONS 15454 Optical Transport Platforms. (These were the high-end DWDM-equipped switches Cisco acquired from Cerent). Cogent has assigned four ONS 15454s to each of its metro rings, which in turn provide bandwidth for up to eight buildings."

"The ONS 15454s shunt the gigabit Ethernet channels from the GSR onto OC-48 SONET links (2.5 Gbit/s apiece). These OC-48 links in turn are sent through bi-directional DWDMs couplers before leaving the POP. The couplers act as passive optical splitters, dividing the light carried by the two OC-48s into 16 wavelengths—eight passing clockwise through the ring, eight passing counter-clockwise. Cogent recently signed an agreement to obtain these couplers from Avanex Corp. <http://www.avanex.com>."

"CPE equipment: In the basement of each

building in Cogent's network, an optical add/drop multiplexer (also from Avanex) picks up two of the 16 wavelengths on the ring—one traveling clockwise, one counter-clockwise. The two wavelengths enter a Cisco Catalyst 2948G-L3 router. One of the wavelengths is converted directly to bandwidth for use by the building's tenants—at 100 Mbit/s per tenant. The other provides an alternate route in case the first link fails due to a break in fiber or a failure in one of the network devices."

"If fiber is cut on the way into the building, for instance, the router will sense that packets aren't coming through its primary port. It automatically will shift to the secondary port. Likewise, if a device fails on the customer network, the Internet connection is still guaranteed."

Please comment further on what the Avanex gear really does for you.

Schaeffer: It allows us to select and recombine the signal so that we can operate WDM service on a single fiber. JDS Uniphase makes similar equipment which combines and splits the signals at purely the optical levels. So in that sense it really is even below SONET. Consider it layer 'zero' for physical split and recombination.

COOK Report: And the Avenex is just better price and performance than the Uniphase?

Schaeffer: Yes. We choose Avenex based on availability and price.

COOK Report: Does it give you some failure protection at the level of each building?

Schaeffer: Not really. We are still getting protection at layer 3 coming out of the routers that we put in each building. Don't confuse Avenex and the mention of SONET with protection. All Avenex is enabling us to do is split and recombine the lightwaves at the physical network layer.

COOK Report: Then Avenex is really a multiplexer?

Schaeffer: Yes. Chromatis is the full system and has the Avenex type of equipment built into it. The Chromatis equipment is in our pops on the large local rings, while our pops on the smaller local rings will contain the Cerent (Cisco ONS 15454 Optical Transport Platforms) equipment.

COOK Report: Do you have SONET framing coming out of the local pops where in each building it is changed to gigabit Ethernet?

Schaeffer: In our local pops we come in to the router from the local fiber ring at OC48

and out of the router as gigabit Ethernet. We then go into a layer three router at the building and out of that for distribution within the building at 100 BaseT.

Fast Track

COOK Report: When did you first come up with the idea for the company? What has been your time line for putting the pieces together?

Schaeffer: I came up with the idea for the company in the summer of 99. In August I hired a small staff. And self funded while we did due diligence to evaluate the technology and availability of the fiber facilities; to survey the market for customer reaction to our product. After about the first three months, in October, we began to seek external funding. We concluded that search in December.

COOK Report: Your external search was carried out with VCs?

Schaeffer: VCs. We raised 26 million dollars from a group of VCs.

COOK Report: According to Light Reading you have raised upwards of a billion dollars. Can you explain how you put such impressive financing together?

Schaeffer: The figures in the article were: Local dark fiber rings from Metromedia Fiber Network: \$100 million Nationwide long-haul fiber from Williams Communications: \$215 million Optical routing and switching equipment from Cisco Systems: \$280 million Multiservice provisioning platforms from Chromatis Networks: \$405 million TOTAL: \$1 billion .

However the figure for Chromatis was hugely mistaken: only ten million. Not \$405 million. The total figures then are on the order of \$600 million. We raised in excess of \$100 million in equity. We do have vendor debt. And we will generate a free cash flow toward that from our operations.

COOK Report: When you sign these deals with vendors in the amounts of tens and even hundreds of millions of dollars, what happens? Do you say we will pay "x" amount over the life of the agreement with "y" amount paid now and "x minus y" paid according to the scheduled dictated by agreement "z"? Is this how it works?

Schaeffer: It could be. Some of our payments are in cash. Some are paid over time and I am probably not at liberty to disclose the terms and conditions of any of those transactions.

COOK Report: I am not surprised to hear that. In other words, however, you can com-

mit over the course of two or three years to the expenditure of say a billion dollars without having to have the full billion in the bank (or in your case closer to \$600, 000, 000) on day one of your operations.

Schaeffer: Correct.

Real Estate Issues

COOK Report: Can you describe for me how you select buildings and what your connect process will be?

Schaeffer: We are selecting large buildings in core central business districts that have multiple tenants. These of course are buildings that are attached to the Metromedia Fiber network. We go into the buildings we have targeted and market our services, not only to the building owners, but also to the tenants where we make sure that we secure their demand. Our service is desirable to the building owner because it enhances the value of his property. However the primary interest is focused among the tenants. What we do that is different from any of the other facilities-based providers. When you secure tenant demand, it is time to go back and negotiate with the landlord.

COOK Report: Have you succeeded in doing just that in a few places so far?

Schaeffer: The answer is yes. Right now we are bound by confidentiality but we are going to be announcing later this fall a series of these agreements.

COOK Report: How will your potential customers begin to find out that you are "for real"?

Schaeffer: We do wave the installation charge for our early adaptors through our charter customers plan. I think the value that we are offering and the credibility we have garnered through some of our capital raised has given us a means of convincing our tenants that we are in fact real.

COOK Report: Oh, I am sure they have an inherent belief, but also that within their organizations you have a bunch of very conservative accounting types who say: 'just to be absolutely sure there are no nasty surprises let's not announce anything until the bits are flowing.

Schaeffer: That's correct.

COOK Report: What are you willing to say (without giving too much away to your potential or actual competitors - I believe that Yipes has been mentioned as potential competitor) about how you target and shape your marketing effort?

Schaeffer: We are targeting our buildings

based on the ability to buy dark fiber to and within them, on the gross size of the building, and the diversity of the tenant base within the building. Once we find a good looking building we talk to the tenants within.

COOK Report: Do you have a sales staff already at work?

Schaeffer: Yes. About 15 very good people. We have several people in some areas and one or two people in other areas. We have chosen people who have usually been in service provider, (telephony carriers or ISPs) management or direct sales.

COOK Report: Someone who had a great track record with UUNET would, I imagine be very attractive.

Schaeffer: Indeed. We have also had a few equipment sales folk join the team.

COOK Report: So what kinds of things do you have to get done before opening for business in the New York market in October.

Schaeffer: We are physically deploying the Cisco optronics and the core routers along our network throughout the country. The opening cities for our network will be New York, Chicago, Philadelphia and Washington DC in October.

COOK Report: What about the remainder of your cities as far as citing equipment is concerned?

Schaeffer: We have all our locations identified in 20 major cities on both the East Coast and the West Coast. These two include the four that will open in October.

Internet Connection and Peering

COOK Report: How are you dealing with the problems of peering and interconnection suffered by other new green field players such as Qwest and Level 3? For example Qwest had to buy EUNet to get necessary peering in Europe and Colorado Supernet for peering purposes in the US. But even ATT bought CERFnet in California to aid in its peering. Certainly, if you say to all the other major players (UUNET, Sprint, Genuity -formerly GTE Internetworking-PSI and so on: "we'd like to buy 100 megabit per-second transit from you at five or six interconnect points on each of your networks," the price quoted would be tremendous.

Schaeffer: We have a team of individuals within the company that is dedicated to a combination of public and private peering.

As we speak they are rapidly entering agreements with many parties. Again those agreements are subject to non disclosures, but let me say that in some cases we are willing to be very flexible in our connections with our partners.

COOK Report: What do you mean by flexible?

Schaeffer: In the case of some very large carriers we are willing to trade transport for peering.

COOK Reports: Fair enough, you do have fiber of your own. What are you going to do at the various interconnect points such as the ones that Equinix is building?

Schaeffer: We are interconnecting at these kinds of exchange points. However I cannot comment specifically whether we will be at the Equinix points.

COOK Report: Talking in general terms when you go into an exchange point are many or most of the tier two and regional ISPs aware of you at this point in time? It might be advantageous to the both of you to just say lets interconnect and do true unpaid peering?

Schaeffer: We are doing this and we were at the most recent NANOG meeting discussing peering with many parties. In addition we are participating at more traditional for profit shows such as ISPCON and InterOp where we are promoting our peering strategy.

COOK Report: One of your people mentioned in a conversation that you were talking with UUNET which has seven conditions that must be met in advance before it will peer. This person stated that you had already met six of the seven - the seventh being the balance of traffic between the two networks. UUNET, of course, is regarded as the toughest peering nut to crack. Any comments?

Schaeffer: We are working on that and are in discussions with those parties. It is our intention to meet all seven of those criteria.

COOK Report: It certainly sounds like you are well aware of the difficulties and complexities behind this issue as you hack your way through the jungle.

Schaeffer: Indeed we are hacking away.

COOK Report: According to Light Reading: "Brad Kummer, former chief technology officer of the optical fiber systems engineering group at Lucent Technologies" is your CTO. Are you willing to say anything about the identity Internet savvy types who

are helping you build this portion of your network?

Schaeffer: We have some of those on staff and have others on the other side with whom we are negotiating. I would rather not disclose individual names.

COOK Report: Are we likely to see anyone on lists like NANOG or IETF between now and October who may say I am from Cogent Communications and I have the following question?

Schaeffer: The answer is yes. Multiple people are devoted to building our role in the Internet community and I expect that by August you will see them doing some of that building even on public lists.

COOK Report: Are issues of your network's handling of IP telephony or streaming multimedia important at this time?

Schaeffer: We expect that the network will handle them just fine. But aside from this these are basically applications and the extent of their use will be up to our customers.

COOK Report: What do you see as the major issues that you need to ride herd on between now and October and the major issues that will determine the degree of your success or failure over the next two years?

Schaeffer: We must execute with precision and rigor against our business plan. It will be about sales, engineering and managing the network.

COOK Report: The fact that you are facilities based should make it easy to increase your own bandwidth on an as needed basis, should it not?

Schaeffer: That is correct.

COOK Report: Talking long haul and metro area facilities based infrastructure, would you agree that Level 3 and Next Link are about the only other players that might be seen as competitors for you?

Schaeffer: Level 3, of course, has plenty of facilities, but don't forget that they are doing lots of circuit switched telephony based things and expensive SONET based things that gives them a very different business model from ours. Because of their need to do CLEC and ILEC interconnects they have a very different architecture for their network from our own which is dedicated entirely to an Ethernet interface for IP over data.

COOK Report: From an analyst's point of view, with their heavy investment in SONET and need of backward compatible PSTN capability, the costs underlying their busi-

ness model are very different and very significantly greater than your own. Correct?

Schaeffer: Absolutely.

COOK Report: Would it be correct to say then that you feel that you are the first one to take the insights about this new technology that Bill St Arnaud posited with CANet3 two years ago and commercially implement them?

Schaeffer: Yes and I think he would agree with that.

COOK Report: I also think he would and the cost effectiveness is down right fascinating.

Cogent Communications, Inc. Gets \$90M In 2nd-Round Venture Capital

July 17, 2000, Dow Jones Newswires

WASHINGTON -- Cogent Communications received \$90 million in second-round venture capital funding led by Oak Investment Partners.

In a press release Monday, Cogent said Ed Glassmeyer, founder of Oak Investment Partners, will join Cogent's board.

In a press release Monday, Cogent said it has received a total of \$396 million in funding, including a \$280 million vendor facility from Cisco Systems Inc. (CSCO).

The company said it will use the money for dark fiber payments, network equipment, hiring more employees and general corporate spending.

Other second-round investors are Broadview Capital Partners, Worldview Technology Partners, Jerusalem Venture Partners, Nassau Capital, Texas Pacific Group unit ACON Venture Partners, Capital Research, Barnard & Co., Comdisco and Boulder Ventures.

Cogent offers high-speed Internet access to businesses.

Company Web site: <http://www.cogentco.com>

-Nick Baker; Dow Jones Newswires; 201-938-5388

GAO Report Shatters ICANN Founding Myths: Neither Privatization Nor Power Sharing Policy Control Over Operation and Content of the "Authoritative" Root - ICANN's Single Biggest Prize - Remains with US Government

The ICANN bandwagon rolls on. The official propaganda says that ICANN's legality was blessed by the GAO and that the Yokohama Board meeting has ended with ICANN's pledge to add new TLDs to the root and at large members to the Board fulfilled. When we look at the details of what actually did happen, we find significant differences between ICANN's allegations and reality. Nevertheless, ICANN is making progress towards cementing its control over the Internet. It has put together a plan to add an undisclosed but small number of new gTLDs to the root by year's end. In doing so it has moved very skillfully to deflect the likelihood of an IO Design lawsuit against it. And at every step of the way it continues to game the process toward outcomes favorable to its IBM, trade mark, and European Commission backers. It operates secure in the knowledge that it has set up such malleable and Byzantine processes that anyone who isn't a full time observer will be unable to distinguish what it is really doing from what it says it is doing. Critical to its ability to game the outcome is ICANN's ability to change at will the legal rules known as by-laws under which it must operate. As one lawyer wrote to BWG: "For an organization that purports to care about the stability of the net, the malleability of its By-laws is ironic. There have now been nine sets of Bylaw changes. (<http://www.icann.org/general/archive-bylaws/bylaws-history.htm>). That's stability?"

The GAO Report

On Friday July 7th the GAO report on ICANN was finally issued. The report www.gao.gov/new.items/og00033r.pdf found ICANN's creation to have been legal. Predictably the naïve general press focused upon this outcome with claims that ICANN was now off the hook as to its legitimacy.

However, the report's real significance was what it said about the Department of Commerce's authority over the DNS Root. Basically it found that it had none. From the report's executive summary we read; "The Department has no specific statutory obligations to manage the domain name system or to control the authoritative root server. It is uncertain whether transferring control would also include transfer of Government property to a private entity. Determining whether there is government property may be difficult. To the extent that transition of the management control to a private entity would involve the trans-

fer of government property, it is unclear if the Department has the requisite authority to effect such a transfer. *Since the department states that it has no plans to transfer the root system, it [the report] has not examined these issues,*" the report's executive summary concludes. This essentially confirms what we wrote in our May issue (published in late March) when we said that ICANN was essentially US government fraud perpetrated to raise confusion over the issue of control over the root.

Milton Mueller wrote in BWG on July 7: "ICANN was presented to domestic (US) constituencies as a privatization. [We see now that] It was not. The government still retains control of the basic assets. The Commerce Dept probably lacks the authority to transfer the assets. ICANN was presented to international constituencies as a sharing of power, a transfer of control from the US govt to a new organization. It was not. It is still in US hands. . . . Two of ICANN's founding myths have been shredded.

And on the politechs list **Mueller** wrote: "by 'privatization' I mean the creation of a clear and definite system of property rights that would pave the way for a competitive marketplace. The US Commerce Department did not do this. It contracted with a private sector corporation to take on some administrative functions related to making policy for DNS."

"That private corporation is becoming more and more governmental because of the power it holds over a unique essential facility (the root of the DNS, and the address space)." [Editor: Or the power it claims to hold.] Read the recent GAO report. The US government retains ultimate authority over the root. The report can't figure out whether the root and address spaces are "property" that can be turned over nor can it say whether Commerce Dept has the authority to turn that property over to anyone. So the government has just contracted out some of its functions. This gives us the worst of both worlds. No procedural protections of government, no competitive checks and balances from the marketplace."

To muddy the issue, at the Yokohama Board meeting when Beckwith Burr was asked to distinguish between 'operational management' of the root and 'ultimate policy control of the root,' and she claimed the GAO report leaves *no question* that the DoC will transfer the former to ICANN upon receiving ICANN's proposal on this point.

Auerbach summed up the situation nicely: A chauffeur has "operational management" of a limousine. But the owner gets to exercise

"policy control" by saying where it goes and who gets to ride in it.

ICANN's First Lawsuit

On June 26 **Afternic.com** sued ICANN for turning down its application for registrar status.

<http://www.nytimes.com/library/tech/00/06/biztech/articles/26doma.html>

<http://www.icann.org/announcements/advisory-20jun00.htm>

Karl Auerbach commented: Hmmm, has ICANN filed an answer to the complaint yet? If not, it strikes me as a particularly stupid tactic for a named defendant to issue a press release that says "we did the complained of act for the following reason." It could well come back to haunt them should they try to later justify their acts on another basis.

BWG member: To the best of my knowledge, there is no restriction on a registrar's ability to also operate as a reseller of domain names.

Auerbach: Indeed, in the material Judith has just posted, it is asserted that NSI's recent auctioning-off of names is essentially a reselling technique.

BWG member: Once again, ICANN 'staff' substitutes its judgment for bottom-up policy. I found the ICANN announcement, linked above, particularly offensive.

Auerbach: And, as I suspect, since "staff" was involved and since "staff" is the corporate legal counsel, their ability to hide behind attorney-client privilege is possibly going to be a bit limited.

Mueller: Any doubts now that ICANN's accreditation scheme is about regulating conduct, as opposed to "technical management"? AfterNIC was denied basically because it was a domain name reseller.

Actually this lawsuit may be more significant for NSI. The registrar agreement contains an explicit prohibition on "warehousing or speculation in domain names by registrars." If AfterNIC is denied on this basis, and I do not think it an unreasonable interpretation of the anti-warehousing provision (even if I disagree with that provision on policy grounds) then on what basis can NSI hold its auctions?

Sounds like discrimination makes for a better case for AfterNIC.

BWG member: Actually, the Agreement just contains a placeholder...

2. To the extent that Consensus Policies are adopted in conformance with Section II.C of this Agreement, the measures permissible under Section II.D.1.b.i shall include, without limitation: ...ii. prohibitions on warehousing of or speculation in domain names by registrars;

To the best of my knowledge, ICANN has not adopted a "Consensus Policy" prohibiting registrars from participating in the warehousing of or speculation in domain names.

ICANN Says "No" and Is Forced to Back Down

COOK Report: Now with wonderful timing Network Solutions, at the end of the preceding week, had announced its own auction. Part of the *Washington Post* article describing the gambit follows:

By Leslie Walker Washington Post Staff Writer Saturday, June 24, 2000; Page E01

"Internet address king Network Solutions Inc. this week notified thousands of owners of online domain names with delinquent registration bills that it would auction the Web addresses to the public unless they pay by Wednesday.

Customers and competitors of the Herndon-based company complained loudly about the plan, saying Network Solutions should instead be required to return all delinquent names to a public database where any of several dozen other accredited registrars could offer for sale the names ending in ".com," ".org" and ".net."

The auction is the latest in what competitors say is a series of anti-competitive actions by Network Solutions, which last year lost its government-sanctioned monopoly on assigning domain names.

"It's a matter of legal interpretation, and our counsel is looking at it," said Michael Roberts, president of Icann, the Internet Corporation for Assigned Names and Numbers, the nonprofit group appointed by the Clinton administration to establish rules for the Internet's addressing system. Icann received the complaints yesterday.

One of the largest and most profitable Internet companies, Network Solutions was acquired by VeriSign Inc. this month in a deal valued at \$15 billion.

Network Solutions said people misinterpreted the e-mail it sent delinquent name-holders to mean that the company would auction names to the highest bidder. "The names will be offered for whatever price is required to liqui-

date the debt," said Douglas Wolford, general manager of the registrar.

Wolford said Network Solutions has to pay a \$6 fee to the registry—the master database maintained by a separate division of the company—even for delinquent accounts. "We have paid for them, and so we are out of pocket," he said."

[Editor: Several paragraphs later the Post article had the following well phrased complaint:] "Elana Broitman, director of policy and public affairs for competitor Register.com, called the auction plan "an abuse of the registry system. They should release the names. This is undermining the competitiveness of the marketplace."

On June 26 in **BWG Karl Auerbach:** I've been thinking about how the lawsuit might proceed... We have the complaint.

ICANN might answer saying: "We are a private corporation, we don't have to do business with anyone, we can be as arbitrary and capricious as we want."

Then **Afternic** might say: "But, but, but, you ICANN are a monopoly and you can't simply act arbitrarily. In addition, you ICANN are administering a government property on behalf of a US government agency and hence are obligated to exercise the same degree of due process as the agency would have."

Then **ICANN** says "Even if we were a monopoly [other roots are possible], we act on behalf of the US government and are hence immune. And even if we weren't, we have made a fair decision in this matter based on our openly and publicly created rules."

Then **Afternic** says: "But your monopoly extends to more than just DNS - it includes the IP address space - and besides you are the vastly dominant entity in the market, not unlike Microsoft. In addition, those so-called rules you mention - they don't even exist, they are spontaneous utterances of an officer and were created in flat violation of your procedures for creating such rules."

Bwg Member: More to the point: In the White Paper, the U.S. gov't. clearly indicated that anti-trust and consumer protection laws would apply to "Newco." Indeed, they are listed as the *only* restraints on Newco as a private entity. This was clearly understood by other governments who participated in the process. Indeed, Industry Canada, in its RFC on the White Paper, expressed concern that this level of protection would prove insufficient.

COOK Report: So with the announcement of its own auction program, Network Solutions, ICANN's cash cow and source of funds without which it would have been unable to pay its bills, created a small problem of equity on the part by ICANN. How could it be possible to deny Afternic accreditation because it sold domain names at auction and not also deny

NSI? No problem. Send in NSI's overseer and have him remind ICANN that some of its charges are more equal than others. On June 28 former IBM lobbyist and current NSI Policy VP Roger Cochetti sent ICANN's Executive Vice President Lois Touton his marching orders:

From: **Cochetti, Roger** Sent: Wednesday, June 28, 2000 4:02 PM To: Louis Touton (E-mail) Subject: CLARIFYING NSI REGISTRAR'S NEW PROGRAM

Louis-

I wanted to clarify the announcement that the NSI Registrar made last week that it will begin making unpaid domain name registrations in .com, .net, and .org available to the general public for limited auction. In designing this program, the NSI Registrar has made every effort to ensure that trademark owners, legitimate registrants, and the Internet community are fully protected; and we believe that it does so.

First, the program will affect only new registrations that are de-activated for non-payment. A registration will be eligible for listing only after the registrant has failed to pay for 90 days, ignored repeated invoices and de-activation notices, and after the use of the domain name on the World Wide Web has been de-activated. Under this program, which affords more opportunity to registrants to pay and avoid service cancellation than most of us could find in other commercial services, it would be essentially impossible for a genuinely-interested registrant not to notice that they were about to lose the domain name that they had registered.

Second, even so, the NSI Registrar will not list or transfer any name that falls under the non-payment provisions described above if it receives an adequate notice from any affected trademark owner that the domain name would violate trademark rights. So, no trademark owner who exercises their responsibility to survey the marketplace for the possibility of infringements should find themselves affected.

Third, to further ensure that there is neither an incentive for, nor even the appearance of, any intentional cyber squatting activity, the NSI Registrar will not transfer any name that is listed in this program for a payment to it that is more than the amount that is due to the NSI Registrar. For anyone who is concerned that the NSI Registrar will make great profit from another's trade name, this demonstrates that our only objective is to recover as much as we can of the amount due to the Registrar.

I hope that these clarifications help put the program in context and I would be happy to discuss them further with you,

Roger Cochetti Senior Vice-President, Policy

Asked to comment a **BWG member** on June

29 said: Basically it doesn't say a damn thing in answer to what I would presume are the same questions Touton posed to Jon Whelan at Afternic:

"... what measures those players (whoever they may be) [puleeze!] taking to ensure that abusively registered domain names are not involved in the auctions? ... what relationships [do] they propose or how [do] they intend to ensure that their accredited status is not abused to disadvantage others?"

Nor quite frankly, should it say anything, because to do so would imply that auction companies and registrars are responsible for policing trademark protection, when in fact, that's the responsibility of trademark owners.

Touton is stuck trying to apply a standard to NSI, that it never should have attempted to apply to Afternic, plain and simple.

Cochetti: So, no trademark owner who exercises their responsibility to survey the marketplace for the possibility of infringements should find themselves affected.

BWG member: ie, "the tm owner would have/should have stepped up to buy or dispute the domain by now anyway ..." [...duh...]

Cochetti: Third, to further ensure that there is neither an incentive for, nor even the appearance of, any intentional cyber squatting activity, the NSI Registrar will not transfer any name that is listed in this program for a payment to it that is more than the amount that is due to the NSI Registrar.

BWG member: (a) What does the amount charged by NSI for the registration have to do with someone buying a domain that's a tm for purposes of selling it to the tm owner or its competitor? Most so-called cybersquatters pay no more than the original \$70 fee now (or less, if they're smart and use competing registrars) ... this is a non-answer that says "it looks like I answered your question so leave me the f*ck alone."

(b) So this isn't an auction, it's simply "who steps in first to pay either the \$35 or \$70" ... ie, the normal first come first serve, which means NSI is simply demanding that these domains be held over for registration by NSI rather than made available as is normally done, for registration by other registrars.

** Looks to me like "their accredited status IS being abused to disadvantage others." ** **Editor:** Of course it is. How quickly everyone forgets that this is all a charade. The agreements that ICANN signed with NSI in the fall of 1999 allow NSI to abrogate it's agreements with ICANN if ICANN's policies cause it financial harm. With NSI there as it's principal source of income, ICANN can not afford to anger it.]

From another **BWG member** on June 29: On

reflection, it occurs to me how badly the Cochetti letter undermines NSI's traditional legal positions:

Traditional NSI position: The registrar & registry bear no responsibility for TM infringement. NSI's attempt to establish a mechanism to assist TM holders after a violation occurs is simply the act of a good corporate citizen.

Cochetti Letter: Registrars (and potentially registries, unclear since we are talking NSI) have an affirmative duty to screen for potential TM conflicts before they occur.

Traditional NSI: Mere registration is not a violation of TM law. Rather "something more is needed" (to quote the Ninth Cir. in *Toepen*). As a corollary to this, the registry/registrar cannot be held liable as a contributory infringer.

Cochetti: Registrars bear a responsibility for policing potential TM violations. Corollary- they may be subject to liability as contributory infringers.

Traditional NSI position: Imposing policing costs on registrars is technically impossible and attempts to do so will increase the cost of registration. As a corollary, registrars cannot be held liable for TM infringement.

Cochetti: Registrars have an affirmative duty to screen for TM infringement. Corollary, registrars can be held liable for TM infringement.

Traditional NSI position: Registration is a service provided by NSI. It creates no actual thing. Once the customer ceases to pay for the service, NSI ceases to maintain the record in its database (rather like a Yellow Pages customer declining to renew its advertisement). Corollary, NSI can alter its contract terms on a going forward basis, and can claim an exclusive right independent of any rights the registrant claims.

Cochetti: A name is a "thing," presumably created by the registrar. When a customer fails to meet its payments, NSI reposes the "thing" and has a right to resell it somehow different from its treatment of inchoate domain names. [Indeed, the logic of the resale raises an interesting question: can NSI pre-emptively register names and "resell" them, if the service it provides is not registration, but manufacture of some "thing."] I cannot even guess at the corollaries that might flow from this one.

Editor: The argument about ICANN accrediting Afternic continued:

Bwg member A: Have you ever had a request for a job interview turned down? Have you ever interviewed for a job and then not been offered the job? Does it matter what the reasons were (aside from the impermissible categories of race etc.)?

Not really. You, me and everyone else has nearly unlimited discretion in deciding who

we want to enter into contracts with and for what reasons - even the Constitution prevents Congress from impairing private contracts. In the commercial context we have anti-trust considerations, but given that ICANN has accredited every single applicant so far except Afternic, it doesn't look as if they've unduly favored a chosen few over the great unwashed. Or, to put it more precisely, a court isn't going to see it that way.

Bwg member B: You are confusing a California corporation with a government agency. What _obligates_ ICANN to accredit anyone? And there would (will, in my opinion) *definitely* be repercussions for ICANN if it started publicly making arguments of this kind. For starters, they could kiss off any money from the ccTLDs.

Bwg member A: Which is precisely why ICANN should accredit Afternic and render the controversy moot. But, there are subtler ways of making the same argument, and any shred of a reason not to accredit Afternic (such as Touton's iBMW.com) is enough ammunition to get the Court to say, "Hey, they could've just exercised near-absolute discretion, but they bent over backwards and found a reason for it. That was fairer than they were required to be." The result is the same, but there's a way to make the argument that doesn't make one look all that bad.

ICANN claims it had a reason. You can yell, "But, Mom, NSI and Register.com are doing it too!" That does not obligate Mom to look favorably upon you. Telling the teacher that the other kids were bad never got you off the hook in grade school, and it doesn't work in court either.

The responsive argument: "ICANN is not perfect. So what."

BWG Member B: two factors restrain ICANN:

- 1) Its bylaws,
- 2) State and federal law.

In this case, the relevant federal law that I was focusing on are the anti-trust laws (i.e., the Sherman and Clayton Acts). California law may also be applicable. ICANN controls a bottleneck facility. This subjects it to more scrutiny for its actions than a mom & pop Internet company.

ICANN's claim that it is a standards body does not get it out of trouble in this regard. Just the opposite. There is plenty of case law that standards organizations and professional organizations can be subject to the anti-trust laws if they are manipulated to function in an anti-competitive fashion.

So afternic.com can prove two things and win: 1) ICANN violated its bylaws in the manner in which it reviewed Afternic's application. Remedy: ICANN must review Afternic's application in accordance with the bylaws (pos-

sibly the court might conclude that ICANN must, under its bylaws, accredit Afternic, but this is a long-shot). 2) ICANN violated the anti-trust laws by applying its rules in an anti-competitive fashion.

Now one, of course, plays into the other. But Afternic could prove (2) separate of (1) or (1) separate from (2). I'm not saying that this is an easy row to hoe. I just maintain that ICANN is not immune from suit, even under the most generous interpretation of the NSI anti-trust suits. I will leave the floor to those who wish to argue that ICANN constitutes some sort of quasi-public body and is subject to due process. From a political point of view, a suit which makes it clear that ICANN has absolute power to discriminate at will (John's position) is something of a win as well.

ICANN Settles Suit

Editor: ICANN never let it get that far. On July 7th we read the following. From a press release just issued by ICANN's PR firm: "Afternic.com Drops Lawsuit, Agrees to ICANN's Terms"

Marina del Rey, CA - (July 7, 2000) The Internet Corporation for Assigned Names and Numbers (ICANN) announced today that <Afternic.com>, Inc., has agreed to voluntarily dismiss its lawsuit against ICANN, and to accept ICANN's terms under which a separate company, eXtraActive, would be accredited.

Under the agreement, <Afternic.com> will:

- * Drop its lawsuit against ICANN;
- * Withdraw <Afternic.com>'s application for registrar accreditation; and

- * Substitute the registrar accreditation application of eXtraActive, an affiliate of <Afternic.com> that will be required to operate as a separate enterprise from <Afternic.com>.

The terms of the agreement are identical to those offered by ICANN to <Afternic.com> on April 7, 2000, approximately two months before <Afternic.com> filed its federal lawsuit against ICANN.

Judith Oppenheimer: Spin observation ... I'm told ICANN released the press release before the agreement was executed.

[The press release claims that:] The terms of the agreement are identical to those offered by ICANN to Afternic.com on April 7, 2000, approximately two months before Afternic.com filed its federal lawsuit against ICANN. . . .

This is rather repugnant ... Afternic was specifically prohibited (see http://www.afternic.com/icann/xa2_60.jpg) from any cross linking, co-branding, etc., with the eXtraActive, the registrar entity; eXtraActive was prohibited from providing registration services to Afternic etc. The prohibitions were

rather draconian, especially in light of the cozy relationship between GreatDomains.com and Register.com.

The prohibitions were specifically reversed (see para F, http://www.afternic.com/icann/touton3_60.jpg and http://www.afternic.com/icann/touton4_60.jpg), and therefore the suit was dropped etc.

Liability Insurance for ICANN Board Members

On Wed, 28 Jun 2000, as the time for At large board Member elections grew near **Karl Auerbach** wrote to the ICANN Board and Staff:

In order for me, or any other person, to make an informed decision about whether to run for Director, several matters must be considered, not the least of which is the extent and risk of personal liability. However, some of the information needed to make that evaluation is not readily available. But it is known to ICANN's management. In order that I, and others, are able to make informed decisions, I would appreciate it if ICANN would answer the following questions by July 10 of this year:

- Is ICANN presently an IRS 501(c) corporation? If so, then under which subsection of 501(c)?

[Certain procedural protections of Director liability hinge on whether the non-profit is a 501(c) and which subsection of 501(c). It is interesting in that the most likely category for ICANN, 501(c)(3), causes some procedural protection to vanish.]

- Does ICANN have a million dollars (US) or more in General Liability insurance?

[Certain statutory liability protections are contingent on the corporation holding at least that much general liability insurance.]

- Does ICANN have a "reasonable" degree of liability protection against the reasonably foreseeable risks?

[Certain statutory liability protections are contingent on the corporation holding insurance that is adequate to protect the public, the exact amounts or means to measure adequacy are not specified by the statutes. Since the measure of reasonability is possibly subjective, it is necessary that the actual policy and policy limits be made available.]

- Does ICANN have an operative Directors Liability insurance policy? If so, then it is important for the potential directors to obtain it as the exact language is critical - general statements that "we have a policy" are essentially useless given the wide variation in coverage afforded by such policies.

- What is ICANN's policy regarding payments to Directors? [The statutes often strip liability

protections if there is any form of compensation beyond strict reimbursement for expenses.]

And on July 6 Auerbach also posted: I got three replies from board members to my second inquiry about liability issues. All were essentially, "staff" will reply. The most complete was from **Linda Wilson** who said "we're busy with Yokohama".

I have requested that, given the gravity of the liability issues and their impact on the decisions of potential candidates, the nomination process be suspended until answers are available. I'm sure that will go flat. But I do hope that board members start to realize that corporate counsel is not their individual counsel and that corporate counsel owes its primary duty to the corporation itself and not to the individual directors. Hopefully a few of 'em go out and get their own advisors and realize that Jones-Day may not be giving advice that is in the best interests of the individual directors.

A **BWG member** commented: I'd be surprised, however, to find that Roberts didn't procure significant D&O liability insurance.

Auerbach replied: I checked the budgets, insurance is lumped under 'administrative' costs, but the total doesn't seem big enough to cover rent, etc and such *and* an expensive insurance policy. (Director's policies aren't cheap - one has to distinguish those from a general liability policy for the corporation.)

An interesting tweak about the "volunteer" protection statutes - at least the California ones - is that they often only protect a volunteer director if the decision complained of was *not* made with reliance on a third party, such as corporate counsel or an accountant. That is a significant reduction in the scope of Director decisions that are protected by those statutes, assuming that all other pre-conditions are met. Roberts is in an "interesting" position, since he is not a "volunteer" and doesn't qualify for any of the statutory protections. And he's been acting *ultra vires*. In more colorful words: his hinie is a blowin' in the wind.

BWG member: but *what* would trigger a situation in which the liability of the boardmembers would become truly decisive?

Auerbach: ICANN is pasting targets all over its body politic. Virtually every UDRP loser is a potential plaintiff, as are all registries and registrars who feel mal-treated. All people who vote in or consider running for board seats are potential plaintiffs. If employees are treated as "arrogant juveniles" there's a lot of litigation potential there too. And we haven't even started in the IP address allocation arena.

And if ICANN does define root system operations requirements and should DNS go awry, then with several billions of dollars per hour accumulating in damages, litigation is likely. Except for blind surgeons who use explosives I think ICANN is one of the country's

prime candidates for the defendant-of-the-day show.

BWG Member: but, then again, I think back to Harold's (?) tribulations in the PETA case... prediction: if board members are named as co-plaintiffs, JDRP will gallop to the rescue.

Auerbach: JDRP only gallops when paid. And JDRP will have conflict of interest between its representation of ICANN-the-corporation and the board members as Directors — you see, if there is liability, its a zero-sum game, the corporation can cut its cost by shifting the payment to the directors, and vice versa. By-the-way, it's not just the liability amount, but the cost of defending that is devastating.

BWG member: Every potential plaintiff who doesn't sue ICANN ends up enhancing ICANN's defense in a later case. I don't like that, but I think it's true.

Auerbach: I don't agree with that. If it were true, tobacco companies would really have achieved god-like immunity by now.

BWG member: Again I agree—in theory. But my moral judgment invariably clouds my tactical judgment. I jumped for joy when I heard about the afternic suit; Berryhill, whose pessimism I have to respect, didn't jump quite so fast or high.

Auerbach: I also thought that the Afternic thing would be quickly settled (before discovery had a chance to start - I suspect ICANN's greatest fear is what could be discovered) with ICANN somehow letting Afternic enter the registrar business. And that's about how it came out, but with the names changed to protect the guilty.

And on July 9th **Auerbach** wote: I just had it confirmed - ICANN does not have IRS 501(c) status. The confirmation was wishy washy, saying that ICANN was formed with the "intent" to obtain such status. But it does not presently have such status.

What I do not know is whether ICANN has had such status rejected or wither their 1023 (the application) is still wandering the hallways of the IRS. (I failed to ask that question.) Does anybody know how long an application usually takes?

And, while an application is pending, is it typical to have an accounting set-aside for taxes that may be liable if the application is rejected? (I doubt that ICANN would have much tax to pay - the degree of its expenses do make the term "retained earnings" somewhat of a foreign concept.)

The Yokohama Board Meeting

Ted Byfield to BWG on July 14: Burr just distinguished between 'operational management' of the root and 'ultimate policy control of the

root,' and said the GAO report leaves *no question* that the DoC will transfer the former to ICANN upon receiving ICANN's proposal on this point.

Byfield on July 15: this ICANN meeting has run amok. The board is arguing that nine isn't a magic number for the Members At Large board members: if, after a six- month study, the board decides the Members At Large election was somehow defective, then they can choose to limit the Members At Large directors to five, or appoint four more by fiat, or do whatever the f*ck they want. One person after another has stepped up to condemn this claim, with the result that a deus ex machina is being transformed into one of these all-too-familiar pseudo-debates that determines the focus of the subsequent discussions.

BWG member: And they will point to the vast comments (all negative) and use them to say that there was substantive input, reaching consensus. Anyone want to bet against that?

Byfield: Nope. What's fascinating to watch is the way in which these discussions unfold. It's EXTREMELY clear that the players (McLaughlin, Dyson, Roberts, Sims, and in this case maybe Abril) are poised to push the discussion in a precise direction. Arguments, references, and even pat phrases ('magic number') are ready at hand in quantity, and you can just watch them wave their magic wands and turn the ground under their interlocutors' to jello. If you sat down with tapes and transcripts of this discussion, it'd be a cinch to prove that this is the case—who first floats a 'possibility,' who isn't involved and how it gets elaborated in the face of solid opposition. Iirc, Touton was notably absent from this maneuver (but it was 3:30 AM [in New York City] and there was a four-alarm fire in the apartment building next door, so I was a little distracted.

And the revisionism involved in *stunning*. The triumvirate-plus was very systematic in presenting itself as the defender of the Members At Large elections against forces opposed to it, and as the authors of a magnificent compromise. This is an EXACT inversion of the truth, and just as assuredly it's a prelude to some mumbojumbo criticism of the phase one Members At Large elections that will result in either a smaller board or appointments for phase two. They *loathe* the idea of user representation, and have opposed it as every step.

It's at moments like that one that I really recognize anew that Gordon [Cook]'s spleen is more than anything else a case of damned good memory. And recognize as well that efforts to 'reform' this organization are quite misguided: They're *expert* rhetoricians, and their talent at tag-team maneuvers is really quite frightening.

Auerbach: The [Members at Large] election *is* already defective. The fact that those running are going to be blocked from their right to obtain the list of electors already destroys

the ability of a candidate or group to coordinate and campaign outside of the manipulative eyes of management and fifth column c-serpents. (I believe that this issue will become *very* ripe once we have the bylaw changes that actually call for an election.)

Jaimie Love: ICANN just voted to create new TLDs. There was surprising support within the board for diversity of models for the test TLDs, for non-commercial TLDs, for different business models, different regions, and even for relatively high numbers compared to expectations by some. Several said it should be 6 to 10, the WG -C recommendation. Others, such as Vint, were very negative on the numbers, suggesting 1 or maybe 2 or 3. The actual numbers for the test TLDs will be determined later, after a review of the applications. It should be noted that Mike Roberts, the ICANN CEO, was very conservative, and is in general quite hostile to civil society type concerns. Esther was very good on the diversity issue, she really pushed this right off the bat, effectively, and Ken Fockler was very good too. Pindar pushed a bit, and even Jon Cohen was pretty positive. Hans Kraaijenbrink was also good on the issues of numbers and diversity. Fitzsimmons and Conrades were negative on the numbers, asking for very slow role out. Linda Wilson jumped into the go slow camp.

One big problem was the \$50k non-refundable application fee, the high end of the staff recommendation. This makes it very costly to apply, particularly given the small chance that a particular application will be approved.

Love continued: One of the issues in the TLD debate concerns the non-refundable fees for a new TLD application. Before the meeting, an ICANN staff paper said the fees would be \$7,500 to \$50,000, assuming 7 to 20 applications would be submitted. During the public comment period, People said the fees would be a problem for the non-commercial proposals. However, the staff proposal, which was approved by the Board today, choose a flat non-refundable \$50,000 fee, the high end.

I have asked Mike Roberts, the ICANN CEO, and some ICANN board members, to raise the issue of whether or not this fee would "disappear" if an application is rejected in the first round, or if the application would have to pay again to be considered in the second round or third round. Mike said he had an open mind, as did one board member, but the issue has not been addressed by the board. We will press ICANN to at least make the fee a credit to another application, and also ask the board to create tiered fees for different types of TLDs, and even fee waivers.

In the first round, there is really very little information regarding what would succeed, in terms of an application, and indeed, the board has indicated that the numbers will be limited, even if there are a large number of "qualified"

and quality applications.

The board also did not make any distinctions between fees for a commercial TLD, that would be worth millions if approved, and a non-commercial TLD that would be run as a public service. Many see this as running the risk that there will be big barriers to some of the types of applications that would meet various civil society goals.

Love concluded: ICANN is giving Becky Burr some award right now — she is leaving DOC shortly. Does anyone know to which firm? Some here say to a firm that represents NSI, but I don't know anything for sure.

.web Is Dead

COOK Report: looking at the Yokohama events, it is quite apparent that IODesign has been chumped and used for canon fodder by ICANN. IO Design acted as though it would know by meetings end whether or not it would be in the root. But true to ICANN form, no decision has been made. While there are a dozen or so expressions of interest on the ICANN web site, with proposals not due until October 1, it is likely that the Yokohama meeting just marks the starting bell for large players who may be taking a much less public approach than IO Design.

IO design, it would seem is damned if it does and damned if it doesn't. I'll can imagine a certain amount of glee on the part of some ICANN folk for whom IO Design has been a pain, setting the fee in part with a gleeful though of its being a mandated contribution to ICANN's ability to continue to restrict entry to the root. If IO Design wants in the root, it has to pay up front with no refund, when it gets turned down - as it certainly will. For there are players waiting to submit applications that ICANN can frame their gTLD in a much more inviting light than .web And there is about 75 days for applications to be finalized. .eu will be one. There will be another from someone with resources well beyond what IO Design can command.

Having not yet been left out of the ICANN Root, IO Design can hardly sue. It will have to wait until Roberts has collected his \$50,000 and informed the .web registry that its proposal just wasn't up to snuff. IO Design has been trumped. Just how effectively will become clear when the applications are opened on October 1. After all, from ICANN's point of view, why should it do anything to enfranchise an old enemy when, by putting .eu in the Root, it can make the European Commission happy? At the same time it can let one or two huge multinationals play with new TLDs that won't have the trademark baggage that .web does? In fact that won't have any trademark baggage at all. Consequently, ICANN puts two new gTLDs in the root and comes up smelling like roses. As with all things ICANN the odds are that the outcome has very likely already been decided. Unfortunately it has figured out a new gTLD strategy that it will make it very diffi-

cult for IO Design to sue, because day-by-day it has been constructing an alibi designed to show that it is not engaged in any kind of discriminatory tactics. And as far as monopolistic behavior over a critical facility - why good sirs: It will have let at least two new players in the door by year's end.

Becky Burr will soon be gone. Like Magaziner disappearing in a cloud of pixie dust back into the private sector. DoC will be ready for the next political appointee to come in and talk about the Internet as the enabler of the little guy while in reality working very hard to ensure it serves only the interests of big business. The Europeans will have been thrown some crumbs and ICANN's operational control of the root will be gussied up to look like policy control. The only thing consistent about DoC's gyrations has been its inconsistency. After all it told everyone a year ago that ICANN would be required to be able to prove in writing all further assertions of consensus. A requirement that has never been enforced. Therefore one would be ill-advised to take DoC's assertion seriously that it has no intention of giving away the root

Appendix One

On November 10, 1999

ICANN and NSI signed a registry agreement that had been mediated by the Department of Commerce. The agreement <<http://www.icann.org/nsi/nsi-registry-agreement-04nov99.htm>> went on to define a consensus policy. 1. A "Consensus Policy" is one adopted by ICANN as follows: "(a) "Consensus Policies" are those adopted based on a consensus among Internet stakeholders represented in the ICANN process, as demonstrated by (1) the adoption of the policy by the ICANN Board of Directors, (2) a recommendation that the policy should be adopted by at least a two-thirds vote of the council of the ICANN Supporting Organization to which the matter is delegated, and (3) a written report and supporting materials (which must include all substantive submissions to the Supporting Organization relating to the proposal) that (i) documents the extent of agreement and disagreement among impacted groups, (ii) documents the outreach process used to seek to achieve adequate representation of the views of groups that are likely to be impacted, and (iii) documents the nature and intensity of reasoned support and opposition to the proposed policy."

Last November most readers of this material failed to note that this document was solely devoted to defining relations between NSI and ICANN and indeed to protecting NSI from ICANN's abuse. For those hoping that the Department of Commerce would place limits on ICANN's capability to ride roughshod over the interests of users it was far too easy to notice the circumscribed intent of this document. DoC was quite willing to protect the financial interests of NSI shareholders and quite unwilling to say anything on behalf of ordinary users. ICANN was free to wave its magic wand and intone consensus when what it was seeing was not consen-

sus but outrage against its dictatorial policies.

Less well recognized were provisions 24 and 25 that essentially show that contrary to its claims of privatization the US government is still the owner of the .com .net and .org franchise. DoC could put ICANN out of business at its pleasure and foreign governments would be quickly shown that the Clinton promise of internet power sharing was fraudulent.

24. Withdrawal of Recognition of ICANN by the Department of Commerce. In the event that, prior to the expiration or termination of this Agreement under Section 14 or 16(B), the United States Department of Commerce withdraws its recognition of ICANN as NewCo under the Statement of Policy pursuant to the procedures set forth in Section 5 of Amendment 1 (dated November 10, 1999) to the Memorandum of Understanding between ICANN and the Department of Commerce, this Agreement shall terminate.

25. Assignment of Registry Assets. NSI may assign and transfer its registry assets in connection with the sale of its registry business only with the approval of the Department of Commerce."

Appendix Two

Editor: From the following ICANN Url <http://www.icann.org/minutes/prelim-report-16jul00.htm> we get the amazingly brazen ICANN statement on By law changes for membership below. Once again ICANN shows its arrogant belief that a group of unelected internet bureaucrats can rewrite the laws by which the organization functions at will. It is as though the Speaker of the House and two staff members decided that they should legislate by administrative order and justify their actions by changing the Constitution and Bill of Rights at will.

Ellen Rony generously provided the following interpretation of part of what is written below. "According to the latest ICANN "reinterpretation" of its own processes, five of the original nine (DOC called them "interim", as does the GAO, but ICANN insisted that they are "initial" board members) must relinquish their roles in November of 2000. The remaining four of the initial board members will be on the ICANN board until November of 2002. This, of course, is preposterous. Four of 19 members unelected and chosen behind the scenes, will be on the ICANN board for FOUR years. And At Large won't get its full representation on the board until November 2002. Thus, individual participation and representation has been dealt yet another blow."

Editor: Furthermore the ICANN junta has decreed that it will spend the next year deciding whether there should be any at large board members at all.

Bylaws on Membership

Whereas, the Board in Cairo adopted a set of resolutions substantially revising the structure and form of the At Large membership and election process, and directing the ICANN staff to prepare and post for public comment conforming amendments to Article II of the Bylaws;

Whereas, the staff's proposed Bylaws amendments were posted for over one month, and the

subject of extensive comment at the 15 July public forum in Yokohama;

Whereas, the Board agrees that certain modifications to the staff's proposal are necessary to clarify the scope and purpose of the post-election study, and to assure the retention of nine At Large Directors on the Board during the entire period from the annual meeting in 2000 to the annual meeting in 2002; it is

Resolved [00.52], that the Board adopts the following amendments to the ICANN Bylaws:

1. The provisions of Article II of the Bylaws are repealed, and replaced with the following:

Article II: Membership

Section 1. General

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

Section 2. Plan for Selection of Five "At Large" Directors in the Year 2000

Five persons shall be nominated and selected by no later than November 1, 2000, to become "At Large" Directors according to a selection plan adopted by the Board. They shall be seated at the conclusion of the Annual Meeting of the Corporation in 2000.

Section 3. Terms of "At Large" Directors

The five "At Large" Directors seated as described in Section 2 shall serve terms that expire at the conclusion of the Annual Meeting of the Corporation in 2002.

Section 4. Temporary Committees

There shall be two temporary committees to assist in implementing the selection plan mentioned in Section 2 of this Article. The first Committee, the Nominating Committee, shall consist of four Directors of the Corporation and three other individuals, one of whom shall be the current Chair of the Internet Architecture Board. It shall have the responsibility for nominating individuals for possible selection as "At Large" Directors. Notwithstanding anything to the contrary in Article III, Section 4, its choice of nominees shall not be subject to review or reconsideration by the Board. The second committee, the Election Committee, shall consist of three Directors of the Corporation and four other individuals, including at least one with significant expertise in monitoring elections to ensure compliance with election procedures. It shall have the responsibility of recommending to the Board procedures for carrying out and overseeing the selection of five "At Large" Directors under the selection plan. Each committee shall serve pursuant to a charter

adopted by the Board, and will cease to exist as soon as it has carried out the functions identified in that charter. Each committee will establish its own rules and procedures, which must be consistent with its charter.

Section 5. Study of "At Large" Membership

Beginning immediately following the conclusion of the Annual Meeting of the Corporation in 2000, the Corporation shall initiate a comprehensive study of the concept, structure and processes relating to an "At Large" membership for the Corporation. The study shall be structured so as to allow and encourage the participation of organizations worldwide, and shall be a "clean sheet" study -- meaning that previous decisions and conclusions regarding an "At Large" membership will be informative but not determinative, and that the study will start with no preconceptions as to a preferred outcome. The study shall include, but not necessarily be limited to, the following issues, taking into account the limited technical and administrative responsibilities of ICANN:

Whether the ICANN Board should include "At Large" Directors;

If so, how many such Directors there should be;

How any such "At Large" Directors should be selected, including consideration of at least the following options: selection by an "At Large" membership; appointment by the existing Board; selection or appointment by some other entity or entities; and any combination of those options;

If selection by an "At Large" membership is to be used, the processes and procedures by which that selection will take place; and

What the appropriate structure, role and functions of an "At Large" membership should be.

The Board shall establish, by the Annual Meeting in 2000, a process and structure for the study that will enable it to meet the following deadlines:

a. The results of the study should be presented to the Board no later than the second quarterly meeting of the Corporation in 2001;

b. The Board shall review the study, and propose for public comment whatever actions it deems appropriate as a result of the study, on a schedule that would permit the Board to take final action on the study no later than the Annual Meeting of the Corporation in 2001; and

c. Any actions taken by the Board as a result of the study that require the selection of any "At Large" Directors should be implemented on a schedule that will allow any new "At Large" Directors to be seated no later than the conclusion of the Annual Meeting of the Corporation in 2002.

2. Article V of the Bylaws is amended as follows:

a. Section 1 is amended to delete the second sentence and replace it with the following language:

"Five of the "At Large" members of the Initial Board, to be determined by the "At Large" mem-

bers of the Initial Board, shall serve until the conclusion of the Annual Meeting of the Corporation in 2000. The remaining four "At Large" members of the Initial Board shall serve until the conclusion of the Annual Meeting of the Corporation in 2002."

b. Section 4(iv) is amended to read as follows:

"(iv) Nine (9) "At Large" members of the Initial Board during their terms of office prescribed in Section 1 of this Article, and any successors; and"

c. Section 6 is amended to delete the words "and the At Large Council" from the second sentence.

d. Section 9(a) is amended to read as follows:

"(a) "At Large" Directors shall be selected pursuant to the provisions of Article II of these Bylaws."

e. Section 9(c)(2). is amended to read as follows:

"2. The term of "At Large" members of the Initial Board shall expire as stated in Section 1 of this Article;"

f. Section 9(c)(3). is amended to read as follows:

"3. The term of the five At Large Directors seated pursuant to Article II, Section 2 of these Bylaws shall expire as stated in Article II, Section 3; and"

g. The last (unnumbered) paragraph of Section 9(c) is amended to read:

"Each Director, including a Director selected to fill a vacancy or selected at a special meeting, shall hold office until expiration of the term for which selected and qualified and until a successor has been selected and qualified or until that Director resigns or is removed in accordance with these bylaws, provided that no "At Large" Director selected pursuant to Article II of these Bylaws shall continue to hold office after the expiration of his or her term even if a successor has not been selected and qualified. No Director may serve more than two (2) terms."

h. Section 11 is amended to delete the third sentence and, in the fifth sentence, the words "or At Large Council" and the words "or more than one At Large Director."

i. Section 12 is amended to delete the second and third sentences, and replace them with the following sentence:

"Any vacancy occurring on the Board of Directors involving an "At Large" Director, whether from the Initial Board or seated under Article II, Section 2 of these Bylaws, shall be filled by a vote of the remaining Directors."

3. Article VII of the Bylaws is amended to add a subsection (c) to Section 1, to read as follows:

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

4. In Article IX, the second sentence is amended by deleting "or At Large Council".

Optical BGP Proposed as Means of Wide Area Interconnection of Bandwidth-Rich Edges Without Burdening Largest Backbones With Additional Traffic Created by Gigabit Ethernet and Dark Fiber Driven Local Bandwidth Explosion

Editor's Introduction: Where customers buy dark or dim fiber for interconnecting LANS, the temptation now is to use only point to point optical links for extending their own direct connections through carrier clouds. By use gigabit or ten gigabit ether net at the customer premise, and by use of optical BGP, customer LAN networks are enabled to cross a WAN. In such a situation, engineering solutions in the absence of layer 2 QoS will focus on simply adding bandwidth. We thank Bill St Arnaud for permission to publish his excerpt which is about 30% of the total text. There is also a Power Point presentation containing many diagrams that is a part of this draft. Just prior to publication we had the following discussion with Bill St. Arnaud.

COOK Report: Will OBGp answer the criticism of those who say there is no way to set up and tear down in short order special purpose commodity bandwidth paths?

St. Arnaud: No. OBGp like BGP setup is very slow.

There are 2 kinds of optical markets - the carrier market which is currently dominating the optical design plane as per your interview with Sycamore and the IP market, for which there has been virtually no development. The carrier market is concerned about fast setup and tear down. That has never been an issue in the IP world because we usually have multiple paths. It is the old connection view of the world versus the Internet packet view. The fight is still not over.

COOK Report: What is the IP market? Why is there no development?

St. Arnaud: I don't honestly know. Part of the reason is so many manufacturers are focusing on the big carrier market. They see that as a lot more lucrative than the traditional ISP market. This in particular includes CISCO. Juniper and Foundry have been the only companies still focused on the ISP Market.

COOK Report: Is my headline accurate?

St. Arnaud: It is accurate, but not the real reason for OBGp. OBGp is intended to allow customers to control and manage their own wavelengths as if it were their own dark fiber without a carrier managed wavelength service.

Optical BGP Networks Discussion Paper

Contributors:

Bill St. Arnaud, CANARIE Inc.,
bill.st.arnaud@canarie.ca

Rene Hatem, CANARIE Inc.,
rene.hatem@canarie.ca

Wade Hong, Carleton University,
xiong@physics.carleton.ca

John Coulter, Cisco Canada Inc.,
jcoulter@cisco.com

Marc Blanchet, Viagenie Inc.,
marc.blanchet@viagenie.qc.ca

marc.blanchet@viagenie.qc.ca
Abdul Abdalla, Bell Nexxia Inc.,
abdul.abdalla@bellnexxia.com

Ian MacDonald, JDS Uniphase Inc.,
ian.macdonald@ca.jdsunph.com

Florent Parent, Viagenie Inc.,
florent.parent@viagenie.qc.ca

Tom Tam, Bell Nexxia Inc.,
thomas.tam@bellnexxia.com

Mike Weir, Globalknowledge Inc.,
mike.weir@globalknowledge.com

DRAFT—DRAFT—DRAFT March 20,
2000 Revised Draft June 5, 2000

Abstract

As customer owned dark fiber networks combined with low cost WDM systems become prevalent, "customer empowered" optical network architectures where the customer controls the optical routing using exterior routing protocols may be possible. OBGp is a proposed extension to BGP for the manipulation of OXCs to permit them to be automatically setup and configured as BGP speaking devices to support multiple direct optical lightpaths between many different autonomous domains. OBGp may also allow customers at the edge to control a subset of lightpaths within another network's wavelength cloud so that they can manage their own lightpath routing within that cloud. With the large number of adjacencies possible using OBGp lightpaths themselves may be used as a direct peering and transit mechanism between consenting ISPs. As a result the big driver for large capacity systems DWDM will may not be bandwidth but wavelength adjacency and transit services. The exchange of lightpaths

may also allow for a simpler mechanism to allow for settlement in peering and transit between ISPs. In future there may even be wavelength commodity markets where ISPs can trade wavelengths and adjacencies on the open market. The proposed protocol extensions may also allow for the deployment of "carrier free" networks, where the customers at the edge control and route lightpaths across a carrier's optical wavelength cloud.

Definitions

In this paper the works wavelength and lightpath have different meanings. A wavelength is a single instance of a wavelength on point to point Dense Wave Division Multiplex (DWDM) or Coarse Wave Division Multiplex (CWDM) system. A lightpath is a contiguous optical path which may be made up of one or more wavelengths connected across optical switches.

Dark fiber is a fiber strand that is provisioned by a carrier, or other supplier, with no equipment attached to the ends of the fiber. It is the customer's responsibility to attach the appropriate equipment at the end of the fiber and "light" it for whatever purpose they so choose.

A dim wavelength is similar to dark fiber, except that the wavelength is provisioned and managed by a carrier. The customer still has the responsibility of attaching their own equipment at the end of wavelength.

Objectives

Traditionally large enterprise customers and smaller Internet Service Providers (ISPs) had only one default path to a much larger downstream ISP or carrier who aggregated their traffic and managed the interface to other networks. With the explosion of customer owned dark fiber and Wave Division Multiplexing (WDM) systems large enterprise customers and ISPs may soon have the choice of multiple paths between themselves and other networks. They network managers will now have the capability to manage their own interconnection directly to other customers and networks independently of an intermediate carrier or service provider.

BGP is the protocol of choice for the management of links between autonomous management domains. BGP management can

become quite complex, particularly when a customer is trying to manage many independent links. As such the automatic setup and configuration of BGP becomes increasingly more important as the number of possible number of direct interconnections increase. In essence that is the main objective of this paper – to propose a mechanism for the automatic configuration and setup of BGP to support many independent and parallel optical links between autonomous Internet domains.

One possible solution is to treat each optical cross connect as a direct path between a pair of BGP speakers. However, this significantly increases the complexity of any single BGP session, particularly where there may be many parallel lightpaths. The alternate solution is to treat each optical cross connect as an independent virtual BGP router with only one input port and one output port. A virtual BGP router can then be set up for each optical cross connect and separate BGP sessions initiated with its peers. This approach is believed to be much more scalable as each virtual BGP router configuration can be easily cloned to other virtual BGP routers.

To date there has been little effort in addressing the requirement to configure, setup and manage wavelengths between domains or to allow enterprises at the edge to manage their own wavelength configuration across a “wavelength cloud”. The conventional solution to date is for a carrier to operate a wavelength cloud and offer a managed lightpath service to the customers at the edge as shown in **Fig 1**. Enterprise networks or smaller ISPs that are connected, at the edge to these optical clouds will generally have limited view into the network and virtually no control of how their lightpath is routed.

A number of mechanisms have been proposed for the management and control of such “wavelength cloud” systems [LRA00]. Most of these systems have been designed on variations of link state interior routing protocols such as OSPF, IS-IS [KRA00] and PNNI [ATM96] or complementary extensions of MPLS such as MPLmS [ARD99]. For complex single domain networks these protocols allow for the optimized configuration and establishment of lightpaths across a single management domain.

Because these networks provide “common carriage” to many downstream customers they require survivable, fast restorable lightpaths. An essential attribute of these networks is the capability to instantiate and route “end to end” optical channels in near real time and to provide capabilities that enhance network survivability [LRA00].

To date there has been little work done on developing protocols for the management of

wavelengths between separate management domains. In addition, customer control and routing of wavelengths across a wavelength cloud is still not possible. This paper attempts to address these issues by proposing extensions to the well know BGP routing protocol for the management and signalling of wavelengths between autonomous domains.

In addition OBGP enterprise customers at the edge will be able to manage their own direct peerings with other networks across a wavelength cloud as shown in **Fig 2.0**. Rather than aggregating traffic the central network cloud allows customers at the edge to manage their own wavelength routing by controlling specifically assigned cross connects on the optical switches within the cloud and thereby extend their own network domains across the central network cloud to directly connect and peer with other network domains attached to the cloud.

There are a number of objectives to this exercise:

1. To provide enterprise customers and ISPs a set of tools for the automatic setup and configuration of multiple BGP sessions to other autonomous domains;
2. To allow the setup of optical cross connects with other peers as a local optimization issue without the prior involvement or notification of the BGP peers in configuration setup;
3. To allow autonomous Internet domains to establish multiple lightpaths between each other with well known quality of service characteristics for the support of advanced services such as voice and video as well as new protocols such as Ipv6;
4. To allow the autonomous Internet domains to automatically transit lightpaths across their network in support of external peering relationships;
5. To provide a mechanism for the management of large single domain wavelength clouds by breaking the larger cloud into smaller clouds and use OBGP techniques to manage a much smaller subset of wavelengths between the clouds;
6. To allow the deployment of low cost terabit routers that use integrated optical switches rather than complex electrical routing engines for the forwarding of packets;
7. To reduce the number of existing multiple network layers and their supporting protocols to one simple universal protocol OBGP for managing the physical and routing layers;
8. To allow edge customers to manage their own optical paths across a lightpath cloud managed by another network entity;
9. To provide a mechanism for distributed Internet Exchange facilities using the exchange and trading of lightpaths between networks to minimize the need for

hierarchical network architectures to interconnect peering networks; and
10. To define a mechanism for the exchange and trading of lightpaths as a commodity on an open market.

Differences Between OBGP and Other Wavelength Management Protocols

It is important to stress that purpose of OBGP is not intended to replace existing proposed mechanisms for managing wavelength clouds. It is seen as an enhancement to those protocols and to provide solutions not addressed by those mechanisms. The following points may clarify the differences:

1. OBGP is not intended as a signaling protocol to support end to end establishment of lightpaths. Rather it is intended for short range establishment of lightpaths to firstly optimize local router traffic flow, and once that requirement has been to secondly allow external peers to establish lightpaths to optimize their traffic flows. The end result may be an end to end lightpath but that is a coincidental benefit and not a primary objective.
2. In municipal and regional dark fiber networks where the individual participating organizations are interconnected by customer owned fiber, multilateral and transit peering is required. Many of these networks are looking at using WDM systems to support direct peers between themselves and other similar minded organizations. Without direct peers a central wavelength management authority would be required to manage the aggregation of traffic and re-routing to the edge. OBGP is intended to minimize the need for any central networking administrator, but instead allow connections to be direct pair wise peerings.
3. OBGP is intended give the customer some reasonable control of the routing of their lightpaths through another entity’s optical wavelength cloud, perhaps as an overlay to an interior wavelength management protocol. For example, a carrier may have a large managed wavelength cloud, but rather than hiding the routing of the wavelengths from the customer, the customer may be given a limited view of the network topology or a choice of possible routes which are subset of all possible routes. In addition, the customer may have optical routes transiting two separate carrier networks and may wish to interconnect its routes through these clouds at some mid point. As a consequence the customer’s ideal optical wavelength topology may be at variance to the ideal optimized topology of the individual carrier networks. OBGP may al-

low the customer's topology to take precedence over the carrier's preferred topology.

4. Large single domain wavelength clouds simply may become unmanageable and too difficult to optimize for traffic engineering purposes as a single domain. The common solution is to break those single domains into many smaller domains which individually can be optimized. Between the domains a more modest optimization mechanism like OBGp could be used.

Background

In the past, for regulatory and other reasons it was very difficult for individual organizations or ISPs to acquire their own dark fiber. More recently however, municipalities, carriers and consortium of various organizations are undertaking open access dark fiber builds in a variety of jurisdictions around the world.

In addition a number of carriers are starting to sell "dim" wavelengths to individual customers and ISPs. ISPs or enterprise customers can now build their own optical network by interconnecting "dim" wavelengths and dark fibers from a number of different carriers. For example, a large enterprise customer may own several strands of dark fiber to interconnect facilities across a city as well as lease dim wavelength to interconnect their campus networks with a distant point of presence which may be half way across the country.

The initial driver for the deployment of these customer owned dark fiber and dim wavelength networks is the significant cost savings over managed services from an incumbent carrier. But a secondary benefit is the development of new applications and services that would not be possible from a traditional carrier managed service.

This model of optical networking is rapidly evolving in the university research network environment in Canada and the USA. Many of the advanced research networks in Canada which are part of the CA*net 3 [BTB00] Optical Internet program are deploying their own dark fiber and wavelength networks.

To date management and configuration of wide area optical networks have been the purview of large carriers. Bandwidth was a scarce and expensive commodity and building wide area networks required special and highly qualified engineering and management skills. But with standardization of Gigabit Ethernet and soon 10Gigabit Ethernet as a simple and effective protocol for carrying traffic in the WAN the management of wide area networks has become relatively trivial. With the recent availability of low cost, long reach lasers Gigabit Ethernet segments can now easily extend 40 to 100 km. Gigabit Ethernet and 10Gigabit Ethernet can also be directly mapped to dim wavelengths on a long haul transparent

optical network.

Now that Gigabit Ethernet can be easily mapped to dark fiber strands or wavelengths it is trivial task for a LAN manager to extend the enterprise network across the WAN. The same LAN management tools and techniques that are used to manage the complex enterprise network can now also be used to manage the relatively simple LAN extensions across the city or eventually across the country. Rather than expecting a carrier to build out a WAN optical network to the customer, customers are electing to build out their LAN to the carrier, (or to the ultimate destination network and bypassing the carrier altogether) via a combination of dark fibers and dim wavelengths. Simply put, the LAN is invading the WAN.

As the result of the availability of low cost dark fiber, inexpensive long reach lasers and standardization on Ethernet as the protocol for wide area and local area networking, enterprise customers are starting to lead the drive for DWDM optical internet architectures. While the immediate impact of DWDM and CWDM technology will be to dramatically increase bandwidth the long benefit of such technology will be to "empower" customer to deploy their own self managed wavelength wide area and long haul networks.

These architectural approaches of "customer empowered" networks may have fundamentally different architecture requirements than that of traditional carriers. For example enterprise customers understand that caching and multi-homing can provide greater reliability than fast restoral and protection on individual optical links. Interconnection and peering to many other enterprise networks may also allow the enterprise or small ISP network to bypass more traditional hierarchical carriers and Internet service providers and establish direct peering with destination ISPs.

Currently, optical networks are primarily used for the interconnection of large network domains such as enterprise networks, ISPs, GigaPOPs and so on. Most of these networks already use external routing protocols such as BGP to manage the interconnection of their respective networks. More importantly it is these large enterprise customers and ISPs who would likely be the first to acquire dark fiber and operate DWDM or CWDM networks. It would seem logical then routing protocols for inter-domain networking might also be useful for interconnecting optical networks.

Advantages of Using BGP for Optical Network Configuration and Routing

Many of the issues that external gateway protocols such as a path vector protocol like BGP

were designed to deal with are similar to the management of multiple wavelengths in an optical network, particularly between management domains. BGP has a basic architecture and tool set is premised on the assumption that it will be primarily used for the establishment of links between independent, autonomous management domains. Although other interior protocols can be used for the management of wavelengths including the management of wavelengths between domains their architectural premise is based on a single management authority i.e. building another network cloud between independent domains.

BGP routing normally only conveys reachability information. It does not convey any information about the optimal topology, quality of service or bandwidth of a particular route. But optical lightpaths (as opposed to SONET channels) are generally of fixed bandwidth – typically 1 and 10 Gbps for CWDM systems or 2.5 and 10 Gbps for DWDM systems. The physical characteristics of a lightpath give it an intrinsic capability of being a "poor man's" logical switched path with a predefined Quality of Service.

Sophisticated intrinsic parameters such as Quality of Service, Restoral or Protection that is commonly available on other types of circuits such as MPLS, Frame Relay, SONET or ATM may be not needed in optical BGP network with DWDM wavelengths between all BGP nodes. An optical DWDM network can be regarded as simply as a more complex inter-domain BGP environment where there exists multiples paths of fixed known bandwidth between network neighbours.

Because a path vector protocol like BGP lists the domains or Autonomous Systems (ASs) that a packet must traverse on an advertised route, the path information enables a customer's router to perform rudimentary traffic engineering on an inter-domain basis. This form of traffic engineering is not as rigorous or complete as MPLS Traffic Engineering. However MPLS-TE, to date only works within a single domain. Inter-domain MPLS will in theory will allow traffic engineered links across domain, but the negotiation and transfer of RSVP and LDP message request across domains will be complex and most likely be mired in non-technical peering and transit business issues.

One of the very useful features of BGP is that it uses TCP (port 179) for all communications between BGP speaking peers. This means that any type of communications channel can be established between BGP speaking peers. The BGP speaking routers or switches do not have to use the actual data forwarding lightpath that for the communication of routing information. The BGP routing information can use any out of band communications channel, including the Internet itself for the communication of BGP routing information between peers.

Because BGP uses TCP for all communica-

Continued on page 26

Executive Summary con't from p 8

particular router, when the ISP needs to move up to something like defining an end-to-end service for a particular customer, there is a big void. Providing such solutions are challenges that happen both within and between providers.

When you start talking about inter-providers, the situation gets worse. Information can't be exchanged. Providers don't want to provide any sort of service level guarantees for infrastructure that they don't control and to which they have no access. Among the things that need to be done are the creation of some sort of common interfaces that providers can use within their own network, but between different systems. For example provisioning and network management itself, and element management. Today, basically most of the network management systems sit in the element level, meaning they tell you what's going on with particular devices in the network. They don't interface well with any of the other business processes within a provider. These are the customer care processes, billing, provisioning, ordering, trouble-ticketing, which is a big problem in providing service to the business-to-business environment. Ironically in this respect the ILECs have better systems than the Internet industry. Caimis intends to solve that shortcoming.

Caimis has two offices. The San Diego headquarters is focusing on the development of the geographic-locations services under Caimis Geo, Inc. The San Diego headquarters will also be doing all of the productization, sales and distribution of all of Caimis monitors and network management systems. The Ann Arbor office will be the leading edge site for senior software and network engineers who are doing customized network management solutions. They will be the ones staying on the cutting edge technology-wise and requirement-wise for the large providers. San Diego will take some of these solutions and develop standardized products that can then be distributed and made available for the smaller ISPs, enterprises, and others.

Cogent Communications pp. 1 - 8,

We interview Dave Schaeffer, CEO of Cogent which in July 17th announced \$90 million dollars in second round venture capital funding. Cogent is the first national ISP to be formed with the business model made possible by the new combination of SONET free fiber, gigabit Ethernet and dense wave division multiplexing. This technology permits IP data networks to be

built for only one one-hundredth of the cost of their voice equivalents less than 10 years ago. It will use a two hundred million dollar 30 year IRU on a pair of fibers from the Williams Communications Network and fiber from Metromedia to connect office buildings in twenty cities nationwide. It will offer tenants of these buildings 100 megabit per second internet connectivity for \$1000 per month. Cogent has signed commitments to spend more than \$600 million on its build out. Without SONET equipment Cogent is protecting its network at layer 3. It describes itself as a facilities-based nationwide switched LAN.

Gigabit Ethernet is used in their local loop. SONET framing in long haul. The SONET framing is a transport standard that is built into the Cisco routers. Consequently there is no additional cost. In the event of a fiber cut instead of 50 milli-second SONET based recovery, they get one second recovery at the router level which, for data, is adequate. Any physical disruption in service would cause the core routers to start re-routing packets. This process takes about one second. Cogent promises to give its customers a full 100 megabits of band width pledging not to put more than 96 on a OC192 - or 9.6 gigabit - loop of 1500 miles in length.

Schaeffer explains how Cogent selected its fiber providers. He notes that: "There are some vendors whose equipment will only work with certain physical shelter spacing. As we evaluated our potential dark fiber vendor, we looked for that network vendor whose shelter spacing would support the maximum number of equipment vendors. Note that the shelter spacings are not inter connect points. They are there only to amplify the signal. All that we need to be able to do is to rent standard size rack space in the shelters." He points out that "fiber can be optimized for the number of wave lengths, for distance or for the power applied to it. There are a number of different design criteria and each vendor takes a slightly different approach to which of those criteria they emphasize with their fiber."

Schaeffer developed a matrix for the evaluation and selection of both his long distance and metropolitan area software. He discusses in detail how Cogent arrived at the decisions for selecting and integrating both sets of equipment. Cogent's Network will open in New York, Philadelphia, Washington and Chicago in October.

ICANN pp. 15 - 21

We describe how ICANN marches onward to the chant of its own version of reality. The GAO report released in July 7 found ICANN's creation to have been legal but it could not find any support for the idea that the Department of Commerce had the right to manage the root. It concluded that it was also very doubtful that the DoC had the legal right to give control of the root to ICANN saying that it did not look exhaustively at this point because the DoC told it that it had no intention of doing so anyway. In Yokohama only a few days later Beckwith Burr proceeded

to muddy the waters by stating that she was prepared to give ICANN operational but not policy control of the root. An ominous development considering ICANN's statement that it was ready to demand that the root sever operators sign contracts with it.

After having caved in rather than face discovery on its first lawsuit, at Yokohama the ICANN Board proceeded quite skillfully to disarm its next, potentially more serious, opponent IO Design. Having given the impression that it might rule on who would get in the root, it collected a number of TLD applicants on its web site, lied that it had consensus for only a very small number of new top level domains, (between 1 and 3) set up a \$50,000 application fee and then announced that it would receive applications August 1 and announce the winners on October 1. In one swift move it deprived IO design of both due process and anti-trust grounds for suing it. It will choose on October first .eu and a gTLD that will have no negative baggage for its trademark rulers. Dot web will not be in the running.

The Board also showed its continuing fear of letting in any outside influence that might expose its operations to public view. For the ninth time in less than two years ICANN changed its bylaws. To ensure continued control by the original Dyson, Roberts, Touton Cabal, it extended once again the terms of four of the original directors through the annual meeting in November 2002. It decided to study the question of whether it should even have an at large membership. In its efforts to make certain that no one outside the original cabal would ever have a voice it entered the following paragraph into its bylaws: "The Corporation shall not have members as defined in the California Nonprofit Public Benefit Corporation Law ("CNPBCL"), notwithstanding the use of the term "Member" in these bylaws, in a selection plan adopted by Board resolution, or in any other action of the Board. Instead, the Corporation shall allow individuals (described in these bylaws as "Members") to participate in the activities of the Corporation as described in this Article II and in a selection plan adopted by Board resolution, and only to the extent set forth in this Article II and in a selection plan adopted by particular router, when that ISP needs to move up to something like defining an end-to-end service for a particular customer, there is a big void. Providing such solutions are challenges that happen both within and between Board resolution."

Optical BGP, pp. 22-24

"OBGP is a proposed extension to BGP for the manipulation of optical cross connects to permit them to be automatically setup and configured as BGP speaking devices to support multiple direct optical lightpaths between many different autonomous domains. OBGP may also allow customers at the edge to control a subset of lightpaths within another network's wavelength cloud so that they can manage their own lightpath routing within that cloud." . . . [We suggest treating] 'each optical cross connect as an inde-

Continued on page 26

pendent virtual BGP router with only one input port and one output port. A virtual BGP router can then be set up for each optical cross connect and separate BGP sessions initiated with its peers." "The physical characteristics of a lightpath give it an intrinsic capability of being a "poor man's" logical switched path with a pre-defined Quality of Service." "The exchange of lightpaths may also allow for a simpler mechanism to allow for settlement in peering and transit between ISPs." "In future there may even be wavelength commodity markets where ISPs can trade wavelengths and adjacencies on the open market."

OBBG Article Continued from page 24

tion BGP switches and routers do not necessarily have to talk directly to each other. In a complex BGP topology all routing update can be forwarded through a central route server sometimes also referred to as a route arbiter or route reflector. The router arbiter eliminates the need to set up a mesh of BGP links. All updates and changes to the BGP routing can be announced to and from the route arbiter. Route servers and arbiters are quite commonly used at major Internet exchange points for this very purpose.

The other advantage of BGP is that it designed

to support unicast routes. The Internet is fundamentally a unidirectional network. Traditional telecommunication networks on the other hand assume that all links are bi-directional and most routing and configuration protocols in use on these network do not distinguish between forward and return paths [CCM98].

BGP is usually configured to work with single point to point connections or one common LAN segment. Optical wavelengths and lightpaths are essentially point to point links. Optical lightpaths, because they are limited by the physical properties of light are unlikely to be as flexible and complex as electronic equivalents such as ATM or SONET circuits. Hence optical lightpaths can be easily be managed by a point to point protocol such as BGP.

A simple, "good enough" and well proven exterior gateway protocol BGP with perhaps some refinements for managing wavelengths and/or lightpaths may be all that is necessary for customer controlled optical networks. As amply demonstrated in the Ethernet LAN world simplicity and low cost always wins over complexity and costly sophisticated solutions.

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

Subscription Rates

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

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

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