The COOK Report on Internet

Akamai Pushes Web Content to Edge

Content Distribution Overlay Independant of Tier 1 Backbones and Peering Is Created Using Proprietary Tools Based on DNS

Avi Freedman as Chief Architect Heads Effort to Place Servers in Networks of 1,000s of ISPs Worldwide via Incentives of Bandwith Saving & Better Response

Editor's Note: While still in college Avi Freedman founded Net Access, Philadelphia's first ISP in 1992. In the ***COOK Report*** basement in 1995, he put in a 56k-connected "POP;" however, in the mean time, he discovered DS3s, backbones, and routing. Consequently the POP never got beyond the 100-pair copper cable from the pole to our roof. Since 1995 Avi has been one of the most widely known advocates for small and mid size ISPs. In 1998 he joined AboveNet as VP of Engineering and in late 1999 he joined Akamai as Vice President and Chief Network Architect. We interviewed Avi on March 17, 2000.

COOK Report: What happened that caused you to leave you VP of Engineering position at Abovenet and go to Akamai?

Freedman: Akamai is a customer of AboveNet. And it was clear that while they had already accomplished a tremendous amount by late fall of last year, they didn't have people who had understood or had worked on big networks and knew how networks break and how they work.

COOK Report: So did you sign on with them in part to get them going in a shift from what the rest of us think of as caching to this intelligent routing that one sees discussed on their website or were they already well under way?

Freedman: They were already well under way. They already had a solution which no one else has, as far as I know. BGP tries to find the closest network by AS path. But that's not really always correct, in fact, it's often not correct. Akamai's solution uses the dynamic performance topology of the Internet to choose the best route. All performance things being equal, the algorithms may also try to push you towards a box, where the data that you want already is. If it is not there, then it will go get it. It can look like caching, but there's a lot of other stuff going on. So that's a diversion, but to answer your question, I went to talk to Akamai and was completely blown away by the number of actually smart, thoughtful people that they have and the openness of the work environment.

COOK Report: Judging by its website, Akamai has a huge network.

Freedman: We have a pretty big network and, although we don't talk about specific size, tens of gigabits per second of capacity can actually be delivered into hundreds of networks.

How Akamai Captures Talent

To answer your first question, I went there, I interviewed and it was like being back in the university. I really enjoy being in environments where there are lots of people who are smarter than I am. It means that you're guaranteed to learn anything you want, if all of you are interested in the same things. And I've been in environments where there are some smart people. But it's perhaps true that there are more smart people here than average.

COOK Report: Are a lot of them MIT?

Freedman: I think there are well over 100. For example the sales engineers at Akamai have computer science degrees from MIT.

COOK Report: What about some of the MIT TCP/IP gurus, like Dave Clark. Are any of them directly involved with the company?

Freedman: No, not directly. That's the interesting thing. Most of the MIT people who got involved were the algorithms people, who really have an important role on their own to play. To the extent people who design the algorithms were with treating networks as black boxes, they had already done a better job than everyone else. Now if you add to that the actual understanding of several specific networks and what goes on in networks overall, you get even more powerful algorithms.

Now I've also never seen a company where all the management is super excellent at what they do. You could view this as a put-down of other companies, but I really haven't seen any other company where every person in management is a so outstanding in their (and sometimes other) fields. I mean, it involves everyone, in every different phase of the company — and, they're completely open to ideas. [**Editor's Note**: At this point, sensing marketing hype I'd begin to get nervous. But I have known Avi **Freedman** since the fall of 1994 and have watched him grow and mature into one of the Internet's more widely known and respected technical figures. From my personal knowledge of Avi I believe his remarks to be absolutely sincere and genuine. Let me say also that I have no financial interests in or ties to Akamai.]

Volume IX, No. 3, June 2000
ISSN 1071 - 6327

On the Inside:

It is a revolutionary concept to me, having spent the last year essentially in California seeing mainly Silicon Valley companies. Whereas, for a company as mature as Akamai today, given that they were doing gigabit per second stuff, and just about to go public, to say to me: "We really need all this networking knowledge in the company. We want to hire you. If you accept, you're hired on the spot because we want you to hire your own staff and we want you to get them on board in the next 4 days" really got my attention.

COOK Report: I can certainly understand why. Help build some background, though, for the reader who doesn't know anything web about caching, the problems of the web as the world wide wait and things like that. Take me how we went through web caching beginning in 1996 and show me how we got from there to Akamai.

Freedman: Okay, I will disclaim that I can give you absolutely authoritative background here, but people like Peter Danzig, who invented web caching, who also work at Akamai, might be better sources. And especially, don't ask me to give exact dates.

COOK Report: Understood.

Origins of Caching

Freedman: In terms of caching, I realized sometime in the '96 range that people were doing some caching. I spent a bunch of time in Australia working with people who did networking and everyone of them did caching, because the cost of getting data there was 19 cents per megabyte, which, by the way, it still is. It is still so costly because even Telstra can't get the quantity of bandwidth, it would like to have. So caching was used for two things. Abroad, caching was used to save bandwidth and therefore to save money. In the U.S., and pretty much only the U.S. performance was the issue.

COOK Report: In other words, move your content out closer to your user, so you don't have to go all the way across submarine cables or satellite to the U.S. to get it.

Freedman: Exactly. And people even did reverse caching, which then looks a bit like what Akamai is doing. In a place like Australia, if they had web content that they wanted available in the US they would send it to a sever located in the U.S. and serve it from there, so it didn't have constantly to be shipped across the Pacific from Australia.

COOK Report: In other words, Australian websites would be placed in the U.S.?

Freedman: In the U.S., yes. Which is somewhat like reverse caching. Which is what Inktomi is selling to people who want to compete with Akamai, but it's not what Akamai does. This gets pretty subtle.

If you look at what Akamai does, it is something that looks much more like a network than like a bunch of caches. In the U.S. in the '96 '97 time frame even, the purpose of caching was to improve performance. So, in fact, you had this box performing the caching and doing it in a way which doesn't always save bandwidth. It actually goes and pre-fetches extra data on the chance that the user might want it.

COOK Report: You might infer that it was actually using more bandwidth than they claimed to save.

Freedman: That's true, some solutions use more bandwidth than they claim to save. But it had a benefit which was that, had you infinite amounts of bandwidth, it would get extra stuff.

COOK Report: From your point of view, extra stuff more quickly. That was the purpose?

Freedman: It would attempt to have the content that the user would want a minute from now in the ready, whether or not it would actualy be asked for. It was a calculated gamble, which is okay if you have a T-3. When bandwidth costs $1,000 a megabit, you're not nearly as concerned as when bandwidth costs $60,000 a megabit. And, remember, until a year ago, bandwidth was still $10,000 a megabit in Europe. It's only now been slowing down into the couple of thousand a megabit or less.

COOK Report: Megabit per second?

Freedman: Right.. In Australia, it's megabyte. But that translates into $60,000 or $80,000 per megabit per second. Let's use Australia to understand both the value of caching and why in some ways they were the world's experts at it for a while. If bandwidth costs you 18 cents per megabyte, how much did disk space cost you?

COOK Report: Presumably, the cost of disk space is going way down.

Freedman: Oh, it's like a couple of pennies per megabyte, if that. So you might as well save everything you've ever seen on the off chance, and if 50%, you know, if anyone ever asks for it again, you win.

COOK Report: But then it gets stale, though, doesn't it?

Freedman: Yes and if bandwidth costs 19 cents a megabyte, you're actually inclined to do really strange things, like tokenize html. Now, "tokenize" is a computer science term. That means that instead of <head>, you'd use the number 1. Instead of </head>, you'd use the number 2. You convert strings of html numbers so that you can store stuff without having to run through a very CPU expensive compression, anyway. In fact, there were people that were stripping banner ads back then, because they didn't want to have to pay to bring them across and ocean to give them to the users.

COOK Report: But if they stripped banner ads in order to save some money, did the advertiser still get the value or was the advertiser bummed out?

Freedman: They simply threw the ads away. And then the other thing was that people would use what I called SECP, Social Engineering Cash Protocol. It's where you tell your users what was and you do it by looking at your cache logs. You tell you users: these are the really fast sites, and the recommended sites for tomorrow. You are confident of their speediness because you know that they are already in your cache. So it's a little bit of social engineering.

But that's the way it happened in many places outside the US.

Now in the U.S., it was a move towards faster connectivity. But how does a content provider react to all this?

A content provider says things like: what this means is we can't get tracking data. Without tracking data we cannot bill our banner advertisers and get good demographics, so we don't like you. That's what they mean.

What they say is copyright. They say, you can't store this; that's a copyright issue. Or they might say: we can't support scaling because of size constraints on number of peers, when what they really mean is, We don't want you to be a better competitor of ours.

There is also the issue for the ISP that, after a certain period of time, the content might become stale. And if you recall, there was a major U.S. backbone that actually implemented hostile caching in their network a couple of years ago and wound up screwing up a bunch of financial transactions.

COOK Report: What kind of caching?

Freedman: I call it hostile. Other people call it different things, like 'transparent'. You are doing it when you intercept users' packets. For example, in Australia, what they do is, they block port 80. Full disclosure. You cannot use our Internet connectivity to get to port 80, you must use our proxy. Full disclosure.

COOK Report: And the proxy goes through

the cache first?

Freedman: The proxy is the cache.

COOK Report: And if what you want is not in the cache, does it let you out port 80?

Freedman: No, it goes gets it for you and gives it you. But you still can't talk on port 80. What the backbone was doing and what caching vendors like Inktomi are encouraging people to do is intercept port 80 packet flows and run them into a box containing their cache without telling the user what's going on. There's a euphemistic term for it, called transparent proxying or caching.. I call it hostile caching.

COOK Report: I see. But if you're doing that, also, and on the website that depends on a lot of my revenue on accurately counting hits and so on, that portion of what my site has seen elsewhere in the world, then I have no idea of what's going on. I can't capture that data. True?

Freedman: Yes.

COOK Report: And I would think people would become quite upset about that.

Freedman: Well, they are, that's why they're Akamai customers. Akamai gives the benefit of caching, it's not quite the same technology, but if you're an Akamai customer, magically, your content is served inside hundreds of locations worldwide, but you get the usage data. You know who's using what images and stuff like that.

COOK Report: The BBN Exodus thing was something entirely different?

Freedman: That was just plain peering politics. The reason that a network cannot compete with Akamai is because their peering politics prevents them from buying connectivity where needed to solve the spot bandwidth problem.

COOK Report: I am hearing about changes that would encourage business-oriented networks to develop along horizontal lines of connectivity and a horizontal mesh overlay of the older, more standard vertical peering practices, where everybody's either a downstream or an upstream customer of somebody else. Content delivery networks is also the newest buzzword for it. It sounds like part of what Akamai is all about.

Freedman: Let's answer that by looking at Akamai's origins. Do you know how Akamai started?

COOK Report: I'm not entirely familiar. My perception is that it started to solve the caching problem or to solve the slow World Wide Web problem.

Freedman: Right. Tim Berners Lee, who invented the Web, came to Tom Layton and Danny Lewin said, you know, we're going to have spot congestion problems.

COOK Report: At MIT surely. Well, the WC3 brought Berners Lee there, so that interaction was probably a byproduct of that, right?

Freedman: Right, and as a result, they started thinking about the problem. But they came about it from an algorithmic perspective, not with a super understanding of how networks connect. If you ask why does the Internet suck, the answer is because packets are being lost between networks. Even though there are networks that can't get packets from one side of the network to another. Most of the loss, a wide percentage, is between networks.

What Akamai Really Does

COOK Report: So in other words, what you've described it for me was the network geek solution. It was created by the network people, people like the Avi **Freedman**'s, who were doing the routing and so on. What you're saying now is that with Berners Lee and the other people at MIT you've been describing, there was a rather different way of going about it that involved algorithmic considerations, and so on. Tell me more about this.

Freedman: For something really authoritative you'd have to talk to Tom or Denny, but I will give you the picture as I understand it. There are a number of different types things that you need to do in order to get to what you could call a content distribution network. You need to be able to solve the problem of, say, ensuring freshness. Akamai has solved the problem. If you use Akamai's network to distribute your content, you will not have stale content being delivered.

COOK Report: Will you not have stale content being delivered because your server will literally be on the network itself? As opposed to sitting in a ISP content hosting web farm somewhere?

Freedman: Well, no, because Akamai doesn't have a network. Akamai doesn't have leased lines to connect its boxes together. Which is good. People say, don't you want to buy a network? And it's like, no. If I bought a network, I'd have to turn it off five times a day, because I have to make sure that I could survive if you turned off. Akamai is immune to fiber cuts.

COOK Report: Help me understand, then, how this is put together in as much detail as you possibly are able to.

Freedman: The base goal is to have servers at the edge of the Internet. What does that mean? It means ideally they're not sitting inside UUNET's hosting center, Abovenet's hosting center, or whomever else's center. Ideally, they're sitting at the ISP's, at the DSL providers, the cable head-ins. Locating in this way minimizes the problems getting to the last mile.

COOK Report: Oh, that's interesting, because for a while the idea was don't have them sitting there at the last mile, get them closer to the backbones, right?

Freedman: Some people might have said that. On the other hand that may not be particularly smart if you're trying to do content distribution.

COOK Report: I heard people saying that one way to get your stuff from one end to the other of the Internet real fast is to bring it into something like an Abovenet content hosting center, where you are very close to the backbone and your ISP has a connection that's close to the backbone, bump, you're on the backbone or OC3 or somewhere else in a millisecond, relatively speaking.

Freedman: Right. But here's another way of looking at Akamai. Akamai is an inverse hosting center, is an inverse ISP. This is the way I described it to people who knew I came from Abovenet. So, at Abovenet, the charter is take the packets and deliver them packets where they want to go with as high quality as possible. Abovenet is still the only company that I'm aware of that does it this way — your customer puts packets into your network by putting a box at an Abovenet hosting center. The way that Abovenet does that is it peers with as many people as it can. It does cold potato routing and pulls the traffic as far along its own network as possible, because its network is, on the whole, less congested than other people's. Now there is still no one else that I'm aware of that has gigabits per second that will go to every little rinky dink peering point in the world and say, hey, I'll peer with you and I'll carry all my traffic from everywhere in the world back and forth to you. No one else does that, really.

However, you still have to go get peering with people, some of whom do and some who don't want to peer with you. You have to wait for telco circuits to come in and wait to get them upgraded as you grow. And in many places, most of the ISPs that actually make up the bulk of the Internet aren't actually at peering points. And a lot of the broadband guys don't want to peer well.

COOK Report: So you are saying, then, that Abovenet does have a business practice where putting your content into its network and bringing your content close to its back-

bone, perhaps does make sense, in comparison at least with the alternatives.

Freedman: Right, but the alternatives there, you know, are a couple of large, well-known commercial networks. Most networks will say Oh, I handed it to them. If they can't get it cross-country, that's their problem. But Abovenet doesn't do that. I was well-trained in what a good network looks like. Now, most networks do not look like good networks. Why? because rumors to the contrary, IP bandwidth is not a commodity.

COOK Report: OK. For example, if you're on those other commercial networks and you hand the packet off to UUNET and the person that wants to receive the packet is downstream on Sprint, one potential problem is how does it bounce from UUNET over to Sprint's network, right?

Freedman: Right, exactly. So that makes sense, that's a good model. But there are tens of thousands of ISP's in the world. Some of them do not go near where Abovenet's fiber goes. And how do you get to those ISP's?

COOK Report: No. How do you get to them?

Freedman: The way Akamai does it is, it puts boxes in hundreds and then thousands of places around the world. For example, there are ways that Akamai could get a good estimate of where all the biggest ISP's were. After this is done my group would combine their traffic utilization data with BGP data and come back with the complete ranked list of how big every network is with respect to Akamai's traffic. There's some very interesting data in there that shows that a number of networks that are not nearly as big as they think are.

COOK Report: What about the MCI Cable and Wireless backbone?

Freedman: The AS3561 network was shrinking before it got divested.

COOK Report: And that AS number was MCI's?

Freedman: Yes. I don't know the reasons but I do know people were having issues getting off that network.

COOK Report: Can you give me perhaps a little more detailed answer as to how you look at how the traffic of people who have their websites on your network, you can trace this traffic and how you monitor it?

Freedman: Networks that are well-routed want to know where their traffic goes, so that they can make sure that at least the large bits of traffic get delivered well. Abovenet is still the only multi-gigabit network that lets you see the state of all their backbone links. Whether it's problems or whether it's good or whatever. So Abovenet is actually a very good partner for Akamai. But that's a different topic.

COOK Report: OK so Akamai does go out and place its boxes all over the place, help me understand exactly what that means.

Freedman: First, networks that are using Junipers and the other big BFR's cannot get data on the new high speed router cards about where the traffic's going.

COOK Report: This is an inherent problem with the design of the router?

Freedman: It's a problem with scalability. And Junipers don't do it at all. So with Cisco, only recently and only on the low speed card, which only goes to four OC3's, can you get the data.

COOK Report: So if you've got an OC48 pipe, you can't get any data.

Freedman: That's right, OC48 cannot tell you where the traffic is going on an OC48 pipe or on a OC12 pipe on a four by OC12 card. What they've said is maybe at some point we'll give you statistical sampled data. But to the best of my knowledge, they don't have that yet. Whereas, one of the best things about looking at Akamai was, oh, here's a multi-gigabit network, in which I know exactly where the traffic's going.

COOK Report: How is that possible?

Freedman: If Akamai didn't know where the traffic was going, then it couldn't bill its customers. In order to be able to bill its customers, it has to know where the traffic is going.

The Akamai Business Model - an Intelligent Virtual Qverlay

COOK Report: So in other words, it's part of my business model. If I want to get my website on the Akamai network, I agree that I'm going to pay by some kind of traffic algorithm of data to and from my website.

Freedman: It looks just like a network. You pay X dollars per megabit on a 95th percentile addressable basis. In order to do that, we need to collect that log data, the billing data from all of our computers about to whom they served how much data. So it turns out that's the same data that tells us, if you mix it up real hard and run it through a little BGP structure of the Internet data to which networks you are sending traffic. It's actually where your traffic wants to go, where demand is. Now, if I don't have connectivity in Joe.net, I may temporarily send that into UUNET or Sprint until I get a box inside Joe.net. So that's the way of looking at it.

COOK Report: Help me understand how you get boxes inside Joe.net. How do servers get where they need to be?

Freedman: I'll just describe to you how it works. So the boxes magically appear in networks, so we start with an Akamai network that has 250 locations. In other words, we have boxes in 250 locations in a couple of hundred networks. They're the networks' locations. We just rent a rack. For a small ISP, we say, if you put our boxes in your network, here's the value proposition. You don't charge us money. We won't charge you for the boxes. The data on the boxes, will always be fresh. You'll even get streams of data, like radio, audio, video, which caches don't do.

If you're broadband, we even have customers who will be able to use our boxes as computers and serve different kinds of content to you. If you have DSL customers in Manila, now we can serve your customers there effectively. If Yahoo wants to have a bigger picture or a stream come to someone who's on broadband, well, we know which of your users are on broadband and which are not.

Now to a backbone we say, put our boxes next to every one of your major routers and if one of your users wants traffic from one of our content providers, they're going to get it from one of our servers. And why not have it be from one of servers sitting in the same city on your network versus elsewhere or versus coming in from a peering pipe? So the value proposition to a big backbone is, again, there's no caching, you don't have to redirect any of your traffic to us. You're not going to have any complaints of staleness, because the content providers are getting the data about who's using what. But you'll magically save bandwidth.

Akamai is a bandwidth-saving proposition to networks. The only people that Akamai charges are our customers, our content providers. So let's say that you sitting at home, have a computer. And you want to get to Yahoo.com, you want to get an image that's off of Yahoo.com. So the Yahoo.com web page says go to something.something.akamai.tech.net or it says something or other that hits Akamai's name servers. Your request for this URL will hit your ISP's name server. In effect if your ISP is Akamaized, you are telling its name server, hey, I want this image, which is identified by this domain name.

Then your ISP's Akamai name server asks a set of eight out of Akamai's total bunch of

top level name servers saying, hey, what is the correct IP address for this domain name? (You can only specify eight at any single time, but there are really more than eight.) And Akamai's name server says, over time, the best place to go is this IP address, which is in such and such a region. Or if it doesn't know, What it says is you'll have to ask someone else.

COOK Report: What would that someone else be?

Freedman: Being another Akamai name server set. So the first set, the first time that your ISP's name server asks the question, it asks it from eight top level name servers that are always the same. And, again, it's eight of the top level name servers, but there are more than eight top level name servers in all.

COOK Report: Are these top level name servers roughly analogous to thirteen root DNS servers?

Freedman: Something like that, except they're distributed on the across different networks and they're supposed to give a good breadth of connectivity quality to far away places. The top level "eight" does not get hit that often. For all the requests to all of Akamai's customers, you're going to have to do that top level query once an hour or so or twice an hour. So it's not hit that often. Because the way the BIND works, it tries all the name servers over time and, over time, it settles on the one that seems to be the best. So then the top level Akamai name server says, Oh, for this customer, you'll have to ask these secondary eight name servers what the actual answer is. Now, of the list of eight that is returned say, six or seven of those eight, are sitting in the same place. They're sitting right in a, let's say, Netaccess or Surf.net in New Jersey or Verio in Philadelphia. Wherever Akamai thinks is a first good stab at close to where you are located. However, notice I said six of the eight. Because two of the eight are sitting somewhere completely different that's completely network topologically different and doesn't tend to go down at the same time.

COOK Report: In other words, in case there's a problem in my locality, I can go somewhere else.

Freedman: For the name servers, we're still at the name servers. Now, when you actually query them, our name servers know the state of the Internet to within a few seconds. And they know whether they should serve you or whether they should refer you, (it's called overflow) to another region. Bumping you to another region might occur because performance is bad or because they're too busy or because a machine has failed or for whatever reason.

COOK Report: In a general sense, can you explain how they know the state of the Internet to within a few seconds?

Freedman: It's a big, distributed computation. That's the general answer. Our name servers each are doing queries about the portion of the Internet that they're assigned to handle normally. Our system knows what users are going to be referred to the second level name servers at any place.

COOK Report: Are these sets of queries that you're set up to do in effect part of your proprietary algorithmic recipes for figuring out what's going on?

Freedman: Yes. It's called the Akamai Secret Sauce. So the top levels get all this data back. Of course, if they don't hear some of the data back, they know that a region is not doing well. Let's say it knows all of a sudden it can't get to Sprint because it's seeing some route flaps over there. Then it will ask who is the second best at this point in time to serve Mr. Cook's name server? And then it'll conclude that it's Joe over there. And then it will say, okay, you need to go to Fast.net in Philadelphia to serve him. But it gives you a real IP address to do it with at that point. The real IP address is returned by your ISP's name server to your browser or actually to your TCP stack. And your TCP stack goes and it goes and gets the content from what has been mapped by this whole transaction, as the server that's going to be fastest to deliver the content to you. And if machines go down, that's okay, if networks go down, that's okay, because there's been all this dynamic performance discovery of which machine will do the job that you need at that particular moment.

COOK Report: And in effect, do I have sort of a hierarchy of service that I can get to, such as if the first one is not available, it goes to the second best one or to the third?

Freedman: Yeah, exactly. We don't say just, oh, the best one is this and only worry about what the best one is with the idea that if you can't get there, we'll just send you to UUNET. This is all actually a pretty complicated thing. This is where the algorithms come in. We don't have a network, but we have eight top level name servers. What happens if they can't talk to each other? Well, there are mathematical ways that you can make a network that partitions well. If you think that you don't need that, you're kidding yourself, because networks get partitioned all the time. Right?

COOK Report: Give me an example of them getting partitioned.

Freedman: Let's say I have a network that has ten cities. One of those cities, the long haul connectivity dies, it's now partitioned. It's partitioned in two pieces.

COOK Report: Okay, a break, a cut.

Freedman: Networks, normal networks, are not designed to be partitioned, so they don't work well. If you design a network that doesn't have an underling transport mechanism that it owns, then it has to be able to work well in the state where it can't talk to each other. It has to be able to say, from my little portion of the world that I can see, this is the right thing to do.

When it can communicate, it sends all this data back and synchronizes everything again with the result that there's a unified network. But in the meantime, all the people that are asking that name server that got partitioned what to do to get a correct answer. All the performance probes that report back, all those calculations are run. It's a really distributed system.

When you wind up getting content from one of our computers, the goal is to make it the right computer. Now, what kind of content can we serve? We can serve html files, we can serve gif images, jpegs. Zip files. We can do streaming. If you want Quicktime, because we run Apple's Quicktime network, or Real Networks or Microsoft streaming media, these are all things that we can do.

We can even start looking at computations at the edge. Doing this gets into how Akamai runs internally inside R&D where there's a big group, which is the group that I work with the most. This group focuses on the way Akamai's used. In one very general sense Akamai is like a big distributed computer that can do lots of things. So one of those things is serving content. Now, does that look like a cache? I don't know, maybe. We can serve streams. Does that look like a cache? Not really. I can serve you a different object at the edge based on your bandwidth speed or IP address or domain name.

Intelligent Traffic Management

COOK Report: Well, are you overlaying kind of a very interesting distributed kind of intelligence system on much of the Internet? A system whose whole purpose is to do intelligent traffic management over the extent of the system that you have?

Freedman: Yes. And part of the infrastructure is given the fact that I have, let's say 50 regions out there. For example, Point A, if it tries to send packets to Point B, is going to get lost through a peering pipe. But Point A could send traffic to Point B and Point B could reflect it to Point C. Getting around the glitchy peering pipe in the process. Let

me give you a real world example but one sanitized to Network A and Network B.

Let's begin with one of our boxes. Let's take an end user and then let's say that there's a network who has pretty bad peering to the world. But note that we have boxes in that network. So once the data gets to that box, we can serve the end users well. But the question is, how do we get our content provider's data to that box?

COOK Report: On the poorly peered network.

Freedman: Right. So if I just try to do a BGP transfer, I'd ride over the Internet. And get packet loss. Which would kill my transfer and make it really tricky. So now I look at where is the content provider centered? Let's see the content provider.

COOK Report: In other words, you may have multiple ways of achieving the goal of getting the packets where they need to be sent and one of the first things you do is look at the various ways that that goal can be achieved.

Freedman: We can do a very precise test to tell us how our regions can get to each other and all our customers, all our content provider customers.

COOK Report: What you're describing to me (I just haven't found the right words to articulate it yet) is that kind of horizontal overlay concept way of thinking about moving traffic. One of the current buzzwords for it is Content Distribution Network.

Freedman: But we are not making networks. I mean, we are sitting on top of the networks. We're not helping networks to peer better with each other, we're just routing around them. So if, let's say the content provider is on Network C. And our user and our boxes, one set of our boxes are on Network A. If I try to go from A to C....

COOK Report: And A is the poorly peered one.

Freedman: Right. So if I try to go from A to C, I'm going to go over a poor peering pipe. And I can't change what that router's going to do. I can't tell that network, oh, to get to the our content provider, don't use C. The network is transparent to me. I can't change what it's going to do. So, but I have, let's say that Network A peers well with five out of fifty of its partners. I have boxes in those five networks and can, instead of going right to C, can reflect my traffic off of B. So I can find the correct, the unfilled, the good quality peering pipes to use to go between networks

It doesn't mean that all traffic from A to C goes well. It just, our content providers' data gets where it needs to go faster.

Because we even work with things that are not static objects, we can do live streaming and things like that. Streaming is not a question of putting the content there once and then serving it many times. You need to build a data stream back.

COOK Report: It's almost a connection-oriented stream, then, isn't it?

"Akamaizing" Content

Freedman: Well, yes, and again, there's no underlying network, because, if I had an underlying network, I'd have to worry about fiber cuts and having it fail. Instead I just rely on having lots of different networks that I can choose from at any given time. We've declared a strategy called Edge Advantage. This means we're going to give people a software interface so that they can run programs on some of our computers that are distributed out in the world. The programs allow our content providers to do, what's called Akamize their traffic and their websites. Which is an automated process. Once they do that, they'll be able to invoke banner ad insertion with just putting a simple single line of html in there principal servers. They can do this because we have banner ad insertion companies. In fact we have every important customer on the Internet. Everyone that, you know, Global Center, Abovenet and Exodus all say that they have 40% of the content. Well, we have all those customers for real.

COOK Report: I'm not quite sure what you mean by that. Are you saying some of your customers also have boxes in web farms at those other three as well as being your customers?

Freedman: All of our customers have boxes in web farms somewhere. And many of our customers have boxes in web farms at Globalcenter, Exodus, Abovenet, UUNET, AT&T, BBN.

COOK Report: I'm still trying to focus on exactly how this works. Is what you're doing, have you developed, it sounds like you've developed a means for putting intelligent devices all over the Internet that talk to each other, that figure out traffic paths and that then can figure the right traffic path to the main web server or server at some big web farm somewhere?

Freedman: That's part of it. But there is an even finer granularity to what we do. Because our content providers tell us, we know what content changes when. We can push content from the content provider to our boxes all over the world and then leave it there till it changes. Or, if the content is dynamic, we can build a communication path back from wherever the user is to the end dynamic content site. Or we can do streams over the Internet, where we reflect them in a sort of zigzag pattern, blasting through congestion to make sure that we have enough redundancy to be able to serve from, say, IDP. We can do many different things.

COOK Report:. What you're saying is, you have this rather large network of name servers but also machines that in effect can function as web servers, data storage devices and so on. Your network knows where your main customer content originating hubs are. Your network figures out how to get people who want the content to those hubs. But your network also then apparently knows how to work with the hubs or the main web farms or websites of your customers and how to bring appropriate parts of that content on an intelligently scheduled regular basis out to your own local disk storage that is scattered all over hell and gone, right?

Freedman: Bingo. You can see why I thought this was interesting.

COOK Report: Indeed. And helped along by the revolution in the crashing costs of local disk storage, too, it makes such strategy, I would think, awfully attractive.

Freedman: Yes, but remember for example, we don't have to push a whole stream. If someone's watching a movie, we don't necessarily have to go get the movie so they can look at it, we can go get chunks of the movie at a time. So, having cheap disk space is cool, but we don't rely on that. We don't have that many super disk farms around.

COOK Report: But the disk farms you do have presumably are highly flexible and highly intelligent and they can have some number of gigabytes or even terabytes of data that probably... I mean, if you look at an average disk farm that brings your customer's data out closer to the periphery, how often does the content of this data change? Very, very frequently, I would guess.

Freedman: Well, it depends. If it's an object that doesn't change, it doesn't change. A lot of the content we serve doesn't change. A lot of the content we serve are live streams, which are one to many, and dynamic. And some of the content we serve is dynamic in nature, like in SSL, so it's one-to-one and dynamic.

COOK Report: Okay, I think Washington Post is Akamaized on your website. You have a big newspaper like Washington Post, are you saying, for example, that they would have their main web servers that they themselves are responsible for getting content changes into and out of, but that you would

take their basic page format and their structural skeletal formats that presumably don't change all that often and store it all over the place.

Freedman: Right.

COOK Report: What do you do with the stuff that changes? If you have stuff on local storage and I'm requesting it through your local storage, there must be some kind of a mating process that goes on between the dynamic, brand-new, fresh data you might have to grab from somewhere else and some of the static data that you might have locally?

Freedman: Let's say you're a content provider, you have a web page. The web page lists the objects and other web pages. It's in the source code for the web page. The object is defined as, let's just say url 1. So then, each region, when our network notices that there's a new url that seems to be popular, it'll be distributed out through our network. Now, as soon as the content provider changes their web page, it'll be distributed, just so long as it's a customer.

Now, let's say that they want to change the image that's at the top of their web page, that effectively looks like it's the top of their web page. They make it a different url. Then it magically is fresh, because people stop asking for the old url. The new url gets magically distributed throughout our network. There's no freshness issue, because when the name of the object changes, we know that we have to redistribute.

That same website may have streaming going on of a live interview. Let's say CNN, because CNN is a customer of ours. We can't push that onto a hard disk because, it's live. So we then set up these virtual communication paths inside our network that enable us to take the data from CNN, distribute it in such a way that it's replicated in to each node. It's not binary. It just adaptively sends extra copies of the data so that there's lossless connectivity all the way to the end user.

COOK Report: Lossless?

Freedman: Lossless. So, that doesn't relate to storage at all. That just gets split and combined and ultimately the clean packets are served, ideally, right at the edge, a few IP hops away from where the end user is sitting because there's nothing we can do with that to make it sit on a hard disk.

COOK Report: How much of this stuff, though, does the customer want, perhaps, cached, so if I miss it, I can come back and get it 30 minutes later?

Freedman: Well, okay, but there's live streaming and there's not live streaming and they're different. But, cached is no longer an appropriate description.

COOK Report: Captured?

Freedman: Captured or put to the edge, whatever. Akamaized. The content provider has to change the name of their things so that they'll fit schema expected by the Akamai name servers.

COOK Report: In other words, if I had a live stream, I could give it one sort of Akamai URLnumber and it would be treated as a live stream. If I wanted that live stream also captured and held on disks for people to get later, I would presumably have some other code that I would assign to the data as well, is that how it works?

Freedman: That's not quite the way it works, because the live streaming versus not-live streaming is a slightly different format, actually.

COOK Report: But do you have a general principle that if I'm a customer and I have data and I want to treat it in various and sundry different ways, you can give me a menu that says here's all the things we can do with your data. The numbers on the menu might run from 1 through 75.

Freedman: We call them type codes and, yes, you can do that. And, yes, that's part of the ARL or Akamai Resource Locator. It's an extension of the URL.

But what I meant to say specifically was, if you're CNN, you can change the objects to not refer to Akamai's name servers and then you're not paying Akamai to distribute that. You have control at any time over whether a portion of your content is being distributed through Akamai.

COOK Report: What you're telling me also is that I can send out, web data, streaming data, whatever, that doesn't get distributed through Akamai, if I choose to do so and that this is data for which I don't have to pay you?

Changing Topology to Ensure Reliability

Freedman: Right. The Akamai system has never failed, but it means that were there a problem, you could just instantly stop using Akamai.

COOK Report: Or in the so-called "template" that I'm just beginning to play around with thinking about, you can stop using the horizontal mesh for distribution and go to the older more traditional vertical stream of the backbone networks and their down streams. In other words, the horizontal way of distributing, the Akamai way of distributing versus the traditional, vertical, send it up from wherever I am into the backbone, let it go across backbones and back down to the edge again. Up, across and down.

Freedman: We'd perhaps describe it a little differently. Let's just say from our perspective, it's a two-dimensional space in terms of where the packets will go. We ripped them out of their plane into a higher plane, which is the Akamai plane, shuffled them around and then drip them back into the what we view as the two-dimensional Internet.

There are other people — like Sandpiper, who is now owned by Digital Island — who actually had similar concepts, although they don't do dynamic stuff as much, and, as a result, they look more like web caching. And you say, what are the differences? The differences are we have more servers in more places and we'll always, and we know exactly where to put our servers. Our networks works, it has never failed, like most of our other competitor's networks have. And the reason is because of our good algorithms.

You know, what algorithms people look at. For example, Tom Leighton, the head of the algorithms department of the math department of MIT, so he's not a dumb person. Some of the questions on which they must focus are, will my algorithm terminate? Will it terminate and give the correct computation? As I scale, how much extra time for it take per unit input? So what's the order of the algorithm? As I scale from 100 nodes to a million, am I going to see the super-exponential breaking of the network?

This is the kind of analysis that yields a network and design that scales and doesn't break, which most of our competitors have had problems with scaling and breaking. This is the head start we have. The fact that networks want us in their network, because we have a few percent of their traffic — you know, we have Yahoo, CNN, Lycos, Motley Fool, all these websites, Barnes & Noble, Amazon — means people really want us in their network, which makes our job easier. When someone, if I wanted to start up a network to compete with Akamai, how would I convince people to put my computers in the pops of their network, give up the colo space and all that, when I have no traffic and therefore I have to pay them for it?

COOK Report: What you're getting at is this is a thing that probably can be done successfully once, period.

Freedman: Right. But of course you never say never. We're certainly not cocky. I mean, someone will come up with the next most brilliant idea.

COOK Report: Or you can do it once and

probably only do it once the way you did it?

Freedman: Right. I've had people ask me, even inside Akamai who are not technical, say, okay, so we respect your opinion, Avi. Is what we're doing really rocket science? Is it really true that we have all these smart things? What I tell them is, Yes, but, at this point mostly we're writing code. The thing is, at any point, we can get 50 people together and have more brain power than anyone else in the network industry to solve any problem or react to any threat or, if the next big thing comes along. I mean, that's the cool thing about working with all these smart people. Because you never want to say no one else can do this, it's just harder.

COOK Report: You feel like you have a better critical mass of talent that you can bring to bear on any given problem.

Freedman: Right. We have 150 super-smart people in R&D. The company grew from R&D, really.

COOK Report: At one level, I see what you're doing, but at another level, I still don't see one thing in it, and I think it's terribly important and maybe you can take me through it step-by-step-by-step. And that is, I mean, you have a network, but you don't have a network. You don't own any leased lines. I think what you've been telling me in part is you've built up a critical mass of customers and tools and servers and so on where you can go to someone like a Cable & Wireless and you can tell them it's to your advantage to give us rack space for our equipment... I mean, what would you tell a big, new network that you're going into and knocking on the door for the first time, what is that you tell them?

The Economics of Distribution

Freedman: We say, right now, it's quite likely that you're paying to take packets from our customers to your customers.

COOK Report: Okay. Because it's the economic argument that I've been missing here and that's what I want you to give me.

Freedman: Or, well, there are two economics. One is the economics of our relationships with our network providers, which is sort of what I've been involuntarily promoted to be in charge of. And the other side is why it makes sense for our customers. Which is very simple, they serve more traffic to their customers, more quality traffic, their customers are happier, they buy more. Or they view more, which lets them charge money for the services they are in business to deliver.

COOK Report: In other words, hey, guy, you're a network, you've got to serve content now, as well as e-mail and mail lists, people are not going to be happy with you if you don't do a good job of serving content. Therefore, we have certain content service solutions to offer you.

Freedman: More importantly, your users are requesting content from Yahoo. There's no question about this. And I can tell you how much traffic per second I'm sending into your network. Even if I'm not sending it directly, I can tell you how much is finding its way into your network. And it's probably finding its way either via transit, which is expensive, or in just two locations, whatever your East Coast and West Coast peering is and then you have to carry it to all its destinations. It's using your long haul capacity. Not only is it using potentially your transit, or your peering, which is limited in the sense that it has to be integrated and costs some money. But it's also using your long haul capacity.

COOK Report: Okay, so you can save me money, how?

Freedman: So if you put my boxes in every one of your POP's, then you have to pay less money. For your backbone is not used as much. You can still charge your customer for delivering the packet to them if that's what your business model is, but you didn't have to pay to bring it from the East Coast to the West Coast to Ohio.

COOK Report: Well, in other words, you'll have to pay once or twice to bring it there, but you won't have to pay to do it 10,000 times a day.

Freedman: Right. And, if you have a problem in your network, we can do the dynamic routing thing, so your customer (for the content they get that's Akamaized) will still get a good experience so they don't get as angry with you if you have a problem.

COOK Report: So your argument is when you come into me and want to put boxes in my POP's, that you're going to give me access to your network and your customers. You're also going to give me access to some of your bandwidth, for which you're not going to charge me. But on the other hand, you get what is a win-win situation, because I can show you how, how in return, for that X percentage of your bandwidth that you're going to give me without charging, you can save Y amounts of money in transit charges and Z additional amounts of money in packets that you only have to get to the edge of your network and onto our boxes one time. Or you only have to get them there once a day instead of a thousand times a day or whatever.

Freedman: Right. But let's look at what happens when you hit a network that doesn't want to work with you? The reason why networks don't want to work with you is because they think they can do what you do. They don't like the fact that when one of their customers Akamaizes, their traffic goes down. Let me just tell you this, I know within a pretty good amount how much traffic every autonomous system on the Internet sends and receives. And who they're connected to.

COOK Report: And you know this because of the number of boxes you already have out there on the 'Net?

Freedman: I know this because I serve gigabits per second of traffic, I know where it goes. I have partnerships that let me get netflow data and cache logs from people, so I know where the eyeballs are. I also, on the flip side, know where the content producers are. How much they are by net, what networks they're on. Let's just say I have a network. Most networks sell to ISP's. The ISP's want us in their network, we go there. Now, they've got T-3's that they're pulling more data than they're sending. Remember a DS3 is full duplex. As much data out as in. In such a situation this network will have some unused out bound bandwidth. So I can go to that ISP and say, Hey, I'll pay you a couple of hundred bucks to get my data out.

COOK Report: I saw you do this a month ago on a mail list. So this is what all your queries were all about?.

Freedman: Yes, exactly. So if a network doesn't want to let us in their network, what's going to happen is, we're going to serve from their down stream customers and they still get no money. Their customer's utilization winds up creeping up or being symmetric, but they still aren't getting any money themselves, so they might as well let us in their network.

COOK Report: And telling them this was essentially just another way of causing people who might be resisting working with you to realize that, if you can get the traffic out and the ISP can get the traffic out, and, if both of you can do it with highest quality and lowest cost, then they too ought to work with you?

Freedman: Yes, but some networks it's true think that they are going to buy some magic boxes, perhaps from Inktomi, and do exactly what we do. Which they can't, because they don't have boxes in other networks. They're restricted to using their own network.

COOK Report: But how all of a sudden did you grow this critical mass so damn quickly and when did you grow it?

Akamai Started Its Deployment from the Big Hosting Centers

Freedman: Yep. Well, one answer is Akamai went into all the big hosting centers. Most of them turned out to be incapable of delivering much traffic, because of congestion, bad design. You know, most hosting centers say, oh, sure, you can burst to 300 megs per second, but when you try bursting at 50, they call you up and say you're hurting their other customers.

COOK Report: So you came in there with a technical solution to a problem that initially they didn't realize they had. But they were realizing it that they had it when you came in.

Freedman: Yes. They were recognizing that, Abovenet was the only place they could burst to a gigabit on demand.

COOK Report: And some of the carrier oriented hosting centers were suffering from a certain Bell-headed mentality that caused them to be slow to understand their environment.

Freedman: And not very interested in sharing that reality with you. In other words, what date am I going to get the circuit installed on? Well, they'll give you a date, even if they know they have no intention of even putting fiber into the building by then.

So the position they were in was, the naive assumption of someone coming from a pure science background is that if someone's in the business of selling bandwidth and they want you to be able to give you the bandwidth so they can deliver it. Makes sense, right? They didn't understand that the Internet came from a model where you sell someone a pipe and hope they use none of it. This mindset sees the network as an expense and everyone wants to minimize expenses. You don't want to over-provision it, because god help you, that might encourage people to use bandwidth. Really, that is the way that many people in the Internet think.

COOK Report: And Abovenet did not think that way. For this is one reason that the web hosting centers, other than Abovenet, would have problems with people who wanted to use serious data in a bursty fashion.

Freedman: Yes. But to give credit where it is due, let me just say that Globalcenter was pretty early on with the gigabit product also. But Abovenet was like, here, take gigabits, because Abovenet realized, if people don't have gigabit speed attachments and they want to do gigabit bandwidth, they can't.

COOK Report: So in other words, some other people began to grok the Abovenet model as well.

Freedman: Right. But groking and changing are two different things. The first generation data centers were simply not of the correct scale. But note also, that AboveNet, who tried to peer with everyone, still can't get into the super edges without years of effort. It's much easier for Akamai to ship a couple of boxes into JoeNet, Omaha's local ISP, than it is to get a telco cross-connect in, or to justify the cost of a ds3 into the 8000 networks who only use a t1 or so of bandwidth.

COOK Report: I infer what you said just by my knowledge of the history of the Internet that the first generation one were those built primarily by the carriers, by the large telephone phone companies. The huge size of the Tier ones did not allow them to move as fast as you would want in the internet age to expand the data centers they had. As a result, the first generation data centers built by the first generation of internet backbone providers are the ones where you're least likely to be able to do a 100 megabits a second.

Freedman: Some reasonable inferences. In this context Akamai is going to be doing peering which for us is a very strange concept. When someone says to me: come to my peering point, I say, what AS's are there and who is likely to peer with me there? The answer to these questions determines whether it's a cost benefit for me to go there. And as far as that goes, I know how much traffic goes to each AS, so I can just do the numbers very easily and say, I need to send this much traffic to, these AS's are there, and they're the ones to which I need to send traffic.

At the same time as Akamai was deploying the big data centers, it started the outreach program that it called the Akamai Accelerated Network Program. By the end of 1999, they already had a hundred of these guys signed up. And by the end of 2000, we think we'll have 2000-3000 ISPs in the Akamai Accelerated Network Program. To put that into context,
the biggest networks may have as many as 300 peering partners.

COOK Report: A hundred of these guys being big data center guys?

Freedman: No, being ISP's. So by the time I got there full time in December, they had two things going. They had buying from networks and data centers and deploying in ISP's. What they didn't have was a set of people that were comfortable laying hands-on routers and such. And capable of talking to networks and making it clear that if they didn't give the bandwidth that they promised, that there would be detection and speedy complaining about said lack of bandwidth delivery.

COOK Report: So when you go into a network and they agree ...

Freedman: To put it bluntly, if I talk to someone and work out the architecture, they know that if they don't do it right, that I have understood exactly what we agreed to and I'm going to go back, trace the cables and get on their case. But unless the ISP's steal our hard disks or plug their boxes into our switches or try to break into our computers, then we cannot get unhappy with the ISP's.

COOK Report: Okay, give me one sentence that describes how you then come into an ISP and what you expect, what you put there. And then contrast that to what you come into a data center with, because you're saying the two are quite different.

How Being Paid for Traffic Changes the Business Model.

Freedman: With an ISP installation, we ship a switch, some computers, a terminal server, consoles, a power supply. Basically, it's all set up, except they have to type the IP addresses into the computers.

COOK Report: To make the configuration mesh into their network.

Freedman: Yes. So, obviously, they have physical access, they can steal the equipment. They could plug their computers into our switch. They could plug our computers into their switch so they can try to sniff the traffic. I mean, we've deployed in countries where they've tried to reverse engineer what we've done. They can take the memory out thinking we wouldn't notice, because each of our boxes ships with a gigabyte of memory, which is worth a few thousand dollars.

So our concerns about ISP's are only that they don't mess around with our gear, basically. We're not expecting, and they're not promising us that they're going to give us 400 megs per second connectivity. Really, they only hurt themselves if they don't give us the best connectivity they can. And they understand that. They want us to be able to steer them traffic to the users, because their users want the data that our content provider customers have.

Now our problem in dealing with hosting centers tends to be that we can't get the bandwidth we want. Because we say, we want to be able to burst 100 megabits a second. We

tell them that we're going to commit to you that we'll do 10 megabits a second. Therefore, you're going to keep available 90 megabits per second. But when we find that we want to use the extra 90 megabits per second, the chances are it isn't there. Why? Because they've aggregated everyone into their routers. Because they don't understand that unless you give people the capacity you said you would for their traffic, you can't bill them for delivering the bits. And people still want to aggregate the hosting centers.

COOK Report: Well, they want to make sure that their traffic capacity in and out is fully used.

Freedman: Yes. That's the problem. Because the way IP works, you don't want to get to where you're fully utilized. You want to get to where you keep your versatility and capacity. You can't engineer to be almost fully used, because then you'll hit congestion. It has to do with the way TCP works.

People are looking at five minute averages, and therefore are not seeing the instantaneous, less than one second bursts that happen quite often. Naive people think that you can keep 90% utilization and that's a good thing to do, because you are using your resources efficiently. In reality, if you're doing 90% utilization, you're squeezing your customer's traffic, which means their flows are smaller, which means they can't send you traffic, which means you can't bill them for it, which means your revenue's lower.

A lot of people are very confused about traffic engineering and what the goal of running a network is. The goal is quantity of service, not Quality of Service.

As Dave Rand (AboveNet CTO) says, Quality of Service is where you say, if you ping me more, I'll drop his packets first. Isn't it better to be in a state where you don't have to drop packets? Where you charge everyone enough money that you can just afford to have essentially infinite bandwidth available for their use? I mean, that's the point, right? That's the point of doing networking is to be able to make happy customers and bill them and all that.

But I haven't answered the question about what is being drafted on Akamai, what network people bring. We're buying OC12's into the biggest networks that don't have good hosting centers.

COOK Report: So you are saying that you are prepared to go to some lengths to get around this problem of people in hosting centers saying: "oh, my god. You want me to be able to hold off a vast percentage of my traffic for you just in case you need it." I gather from what you said, they're rather resistant to that.

Freedman: Well, the business people agreed to it, but their technical people refused to implement it.

COOK Report: So are you saying, then, that in order to deliver the kind of Quality of Service you want to and the kind of service that gets you paid for the bits you deliver, your business model now is changing such that you are going to begin or you actually have begun to buy some of your own bit pipes?

Freedman: Yes. We are now beginning to manage routers, because that's, not because we want to, but because that's the only way we can get into some of these large networks, because you can't get in through their hosting centers, because their hosting centers are not scalable. So, fine, we'll make pipe connectivity to you. No problem.

COOK Report: And you've got to attach those big pipes of connectivity somewhere else.

Freedman: We'll just rent some space, find the closest place where the OC12 can terminate into their network, get some space and stick some boxes right next to it.

COOK Report: But now for the first time you're going to start to build and physically run your own network.

Freedman: Ah, but it's not a physical network. The point of it is, to get the OC12 in, we're not aggregating that OC12. Nor are we aggregating the network A's OC12 with the network B's OC12's and making mega-POP's. We want as fine an overlay mesh as possible. We don't want more than a percent of our connectivity going through a single point of failure.

COOK Report: I understand that. But what if you've got, say, on the face of the Earth, to choose an arbitrary number, five data centers that are giving you problems of being able to get the right connectivity into them. Let's say you're thinking of an OC3 into each of these five data centers. OK, you've got five places on the face of the Earth that your OC3's are going to begin, where are they going to be terminated?

Freedman: In the hosting center. We'll stay in the hosting center. We'll just get an OC12 into their backbone, instead of riding in their hosting center aggregation architecture. And if they don't want us in their data center doing OC12, we'll get a little hut in the back or closet somewhere and, because if the OC12 falls offline every now and then, it's no big deal, since we have more capacity to that network than we had before.

COOK Report: So you'll wind up having to pay some kind of X quantity of money to each of these five hypothetical centers that you wouldn't otherwise be paying for, but in return for the benefits that you see getting out of it, it's a worthwhile expense.

Freedman: Well, that's the other stupid thing, is that people charge more in their hosting centers than they do for their backbone connectivity.

COOK Report: So the amount of money is non-trivial?

Freedman: I'm just saying it doesn't wind up necessarily costing us more money to get a burst of OC12s than to try to get into the data centers. Because the dollars per megabit is lower on OC12s than it would be were I inside big peoples data centers.

COOK Report: Give me a specific situation.

Freedman: Joe Very Large Network wants to charge me $550 per megabit inside their hosting center. That's their bottom line price because their backbone charges them $500 a megabit, because they want to make profit centers. But they'll sell to an ISP at $400 per megabit.

COOK Report: Why do they do that?

Freedman: They try to sell their hosting center to content providers as a premium space and power and one, theoretically, some redundancy and scalability of facilities. But that's a different question.

COOK Report: There are extra charges for the physical facilities, but they make the bandwidth a little bit cheaper?

Freedman: No, they make the bandwidth more expensive at hosting centers. They charge $500, $600 a megabit per second in the hosting centers. But remember that these older hosting centers are each connected by only two OC3's and they have a hundred customers in them. Now how could they possibly give Akamai 300 megabits per second in connectivity if they only have 300 megabits per second into their hosting centers which have a different AS number than their network. That's how seriously these are different networks.

An ISP is not going to be buying into a hosting center. An ISP is going to get a T-3 or an OC3 into a big carrier's router. So we say, we're congested. They say, oh. Okay, we could in four months give you a port on this router rather than running through a switch. Well, but what router? Oh, the customer attached router, okay put us on the backbone router. Okay, we'll put you on the backbone router. Where's the router? Oh, the backbone router still only has an OC3 into the big carrier network. And still has a hundred megs per second of traffic on it.

They can't provision additional OC3's, because they're not doing a gigabit architecture. The other carriers are all the same way because they are all using Cisco 7500's and a couple of OC3's. But the irony of it is they're charging more, because this guy who wants three megabits per second of connectivity, it's worth it to him to pay 50% more for the bandwidth to not have to worry about running a router. But for us, that's unacceptable.

COOK Report: And the ISP hopes he gets adequate service. Or does he really get adequate service?

Limits to Scalability of the Vertical Peering and Backbone Connectivity Model

Freedman: But he does get adequate service. Because if you have fifty people using 3 megabits per second, most of them don't start using 3 megs per second the next day. Most of them grow very slowly. That's how the carrier makes the model work for the ISPs. But for the big websites huge need is there right now. Therefore, they can't be accommodated in this architecture.

So what does that mean? That means that Akamai was already on the path towards moving out to the edge. Broadband, big ISP's, like that. But some of the, but how do you get into the bigger networks? The answer is not through the hosting centers. The answer is now Akamai needs to start looking like a network — even though it isn't... Now this is where it can hurt your brain.

COOK Report: Yeah, right, go ahead.

Freedman: So we'll get OC12's. Now there are some networks that don't want to put us in their network, but have completely open peering policies.

COOK Report: But you get an OC12, you get this damn OC12 from exactly where to exactly where? You get it from the web farm around to the backbone such that you've got a separate link that may be 100 feet or something?

Freedman: It might be sitting inside Joe's Network, where our rack has more connectivity than the whole data center does, because that's what we need. And the leased line guys, leased lines don't get aggregated as much. And by buying OC12's, you make sure, I guarantee you, if somebody's aggregating an OC12, they're under congested in a lot of their backbone. Most people don't have customer-attached routers you can put OC12's on. Or, if their hosting center is full or they see what we're doing and they don't like it, then we'll rent a room somewhere, put the OC12 in or we'll bring it into another data center. Or whatever, but we will do it in such a way that we begin to aggregate multiple sites, but not grow a long haul network.And I don't want, if I'm sitting at Abovenet, and I have a gigabit to them that I can actually use, I don't want to start bringing OC12's in there, because that means, if Abovenet has a power problem — they never have, but let's say they do — a lot of my connectivity is affected. I don't want to aggregate these things and put them near each other. So it might be inside the same hosting center, it might be somewhere else, whatever, we get an OC12 into their router.

COOK Report: Well, I have one other parallel question. Does the structure that you have provide a means of helping big websites deal with distributed denial of service attacks? Or is that totally orthogonal?

Freedman: Yes, it does. But wait: are you talking networks or customers? Networks or content providers.

COOK Report: Content providers.

Freedman: Oh, yes, sure. Because if someone attacks one of our regions, the region just falls offline. Our network operation center doesn't, isn't required, really. Because the network self-heals.

Now, we don't go out there saying, buy our product and you'll never be down. Buy our product and you're immune. But if someone attacks, let's say they're attacking content provider A and one of our regions in a hosting center is right next to that, that region may go offline. But our algorithms automatically redirects the requests somewhere else. Remember when I told you the mapping, some of the name servers handling each region are actually elsewhere. So if a big hosting center in Bay Area goes offline because of a denial of service attack, it doesn't effect anything in our network and to the extent that the content was pushed onto our network previously, then we still have it.

COOK Report: If back in February, Yahoo was seen from several parts of the U.S., if I remember correctly, to be essentially unavailable. But from other parts in the world, Canada and I think abroad, it was quite available. How come?

Freedman: This is actually a pretty long answer and there's another product that we have that really helps with this, which I'm not sure I can tell you about. But this is a complicated answer to get correct, because Yahoo is a customer and they do buy a few different services from us. But the bottom line is if someone is only serving html from a couple of places in the world, then it has nothing to do with Akamai. Unless they're serving their front page through us, if they get attacked, they're not going to be available, they're going to be offline. But there is a lot that relates to what we do and how we do that provides good protection against denial of service attacks.

COOK Report: So where is all this headed?

Freedman: Well, my group at Akamai is integrated. We have two main functions. We work with R&D, especially in the realms of how the computer is attached to the infrastructure, because that's sort of controlled by R&D, so moving from fast DSL to Gig Ethernet and beyond, that's a network architecture issue. And my group is network architecture. We work with the mapping group to better understand how the Internet fails.

COOK Report: The mapping group being the group that looks at traffic patterns?

Freedman: Yes, there's a group that builds a map that's supposed to be accurate to within seconds of the complete performance reachability information of the Internet. I wish I could spend all my time in that group, but my people are needed to help with the building of the network also, because that's really interesting stuff.

Thus we work with the mapping and performance analysis groups in R&D, but we are, for now, really directing the network architecture and strategy and deployment parts of the company to ensure three things. We want our networks to be burstable, with excellent performance and to also be the most effective at the lowest cost.

COOK Report: But your network's not really a network, it's kind of a virtual overlay.

Freedman: Right. Our virtual overlay network. Our bunches of boxes.

COOK Report: And where you are needing them, you are getting either semi-permanent or temporary patch solutions within a data center to get from point A to point B through the data center's network or if the network is performing poorly, you're in there to negotiate, how you will go from point A to point B, via detour of this separate OC3 cable if need be.

Freedman: Right, or into the backbone. But we're also deploying peering regions. Because there are bunches of providers out there. And there are some people think, like from your question, who think that we're caching which is OK. What they don't realize is, if they just BGP peer with us, data magically appears. Consequently, we're going to set up a router. We're going to peer-

ing points. We're going to be peering with people over peering points and begin looking sort of like a network from that perspective.

COOK Report: So you're extending your virtual network, you're extending it a bunch of different and new ways.

Freedman: Yes. And in Europe and Asia, bandwidth isn't necessarily costly, necessarily, if I do it right, because let's say you're a network that's peered in Asia. And I come to you and say, I can give you ten megabits per second to give to all your peering partners and I'm not going to charge you for it. What's the value proposition? Those peering partners would otherwise have to pay to bring that data in over international circuits. That's very valuable data. I call our boxes sometimes Magic Traffic Generators. I can put them there and make more traffic appear, as long as the quality's good.

So it helps them. When you're in their network, if you can double bandwidth at no cost to you, you all of a sudden look more attractive to peers internationally, because you're giving them much more traffic. Therefore it helps us to widely deploy in peering points there and even if we don't peer with people, just by physically putting our content there, they may decide to give us their good connectivity.

Now in the U.S., let's there's a network A, that doesn't want us to deploy inside their network. But they peer with 50 other networks. They don't want to peer with us because they think they want to compete with us.

But they peer with all these people. Therefore all I have to do is go to the people they peer with. Now let's say you, Gordon, are one of these networks that peers with network A. And I say to you, okay, you're cost is $25 per megabit for these circuits. Let's say I put boxes next to each one of your peering pipes and pay you $100 per megabit. Now you're sending more traffic to that network, so you look more attractive as a peer and you're making some money and you're making three times what your cost is.

That's another way we can deploy. So it's just using creativity to figure out what the best way is to get our customer's traffic to the end users of networks who, in some cases, aren't yet on the page in terms of letting us deploy all over them.

COOK Report: And so obviously they've come to you as someone who, over the last several years, has developed a very, very good knowledge of the players and how it works; who the people are; the technical scalability issues. Plus, if somebody tells you they're going to do such-and-such at a big data center, you know and they know that you know they know, that you can check up on them and hold them to what they've agreed to do.

Freedman: Yes, that about wraps it up.

Some Akamai Traffic Statistics

[**Editor's Note**: An email discussion with Avi Freedman on April 15 - 16 yielded the following postscript..**]**

Since Akamai bills its content providing web sites in terms of the aggregate content requested from each of the sites world wide, it must keep good track of the aggregate data shipped to its servers at the edges of the Internet. Or, in more traditional e-commerce terms, data delivered to end user "eyeballs."

When Akamai aggregates this data, it does so in terms of the AS (Autonomous System Number) of each network. While there are a few AS numbers that would generally not be thought of in terms of ISPs numbers and a few ISPs that have more than one AS number, ISPs globally tend to have a single AS number for routing purposes. While the size of the Internet can be measured in more than one way - for example by numbers of email accounts - when networks are ranked according to their percentage of Akamai traffic received one can certainly get some idea of their relative size. The top four globally in terms of percentage are UUNET, Sprint, AOL and ATT. The total for ATT here is combined ASNs of ATT Worldnet, Cerfnet, and the old IBM Global Network (now called AGNS). AOL's > 10% is the combination of a number of AS numbers.

Note again that the following networks are listed in terms of where the eyeballs (Akamaized bits are being sent) are, not from where the bits are originating. Only ISPs and other networks that are multi-homed; truly Tier 1; or got an ASN before 1996 or so will appear in this list. Finally one must note that the table counts all traffic that CAN cross any given network. So let's say that on its way to Joe'sNet which terminates 0.20 of Akamai's total traffic, the traffic transits UUNET or Sprint or any other backbone. In such a case the traffic in question is counted both for the transit backbone and for the net where it terminates. The actual totals of the percentages add up to roughly 277% of aggregate traffic meaning that the average IP prefix is multi-homed 2.77 times.

There are 2 networks > 15.00 %,
UUNET and Sprint

There are 2 networks > 10.00 %,
AOL, and ATT

There are 4 networks > 4.00 %,
There are 6 networks > 2.00 %,
There are 4 networks > 1.00 %,
There are 34 networks > .50 %
There are 56 networks > .25 %
There are 22 networks > .20 %
There are 44 networks > .15 %
There are 85 networks > .10 %
There are 61 networks > .075 %
There are 139 networks > .05 %
There are 77 networks > .04 %
There are 95 networks > .03 %
There are 197 networks > .02 %
There are 487 networks > .01 %
There are 196 networks > .0075 %
There are 378 networks > .0050 %
There are 721 networks > .0025 %
There are 885 networks > .001 %
There are 2849 more networks

Traffic to Europe is accounted under each large transit AS. Europe is about 8%, Asia about 3%, South America about .5-1%.

Listed in alphabetical order the 14 networks in the greater than 1 percent but less that 10 percent categories are Abovenet, BBN, Bell Canada, Concentric, Digex, Exodus, Globalcenter, @Home, Level3, PSI, Qwest, Splitrock, Teleglobe, and Verio.

WIPO Dispute Resolution Panels Creating Their Own Law

ICANN has always been touted by its backers as an opportunity for noble application of the principals of industry self regulation to the Internet.What this means in practice is that industry is able to ignore democratic and procedural safeguards built into the regulatory systems of democratic nations.and act in an arbitrary and high handed manner to accomplish the goals of its powerful supporters. We are now seeing the domain name dispute resolution panels acting in the same way - flaunting their independence from centuries of western juriprudence. Consider the case of eResolution.com http://arbiter.wipo.int/domains/decisions/html/d2000-0110.html

It proudly proclaimed on April 10, 2000. "Although entitled to consider principles of law deemed applicable, the Panel finds it unnecessary to do so in any depth. The jurisprudence which is being rapidly developed by a wide variety of Panelists worldwide under the ICANN Policy provides a fruitful source of precedent. "

Commoditizing Bandwith - Part 2

Bringing Risk Management to the Internet

Lin Franks Explains Issues Involved in Commoditzation

Describes Her Evangelist Role in Management and Training Issues Inherent in Trading Bandwidth

Editor's Note: Lin **Franks** is a Senior Manager in Andersen Consulting's Energy Practice where she helps clients develop points of view and strategies around the converging communications and energy markets. She is a contributing author to The US Power Market, a 1997 book published by Risk Publications, and the forthcoming publication, The Telecommunications Revolution. Her previous position was Director Commodity Market Development, Enron Communications. We interviewed Lin on March 15, 2000.

COOK Report: Please begin to explain more precisely what you're doing and the role that you're performing.

Franks: My unique area of expertise is in the development of emerging commodities markets. My background and my history predates, for instance, the New York Mercantile West Texas Intermediate crude oil futures contract. I've been involved in or have experienced at some level the "commoditization" of several products: crude oil, natural gas, electricity and now bandwidth. If you look at the process that a market or an industry goes through as they are deregulating or undergoing some other significant change, you will find that there are some consistencies from industry to industry. The critical knowledge I bring to bear on each new emerging market is knowledge obtained from having been through the earlier evolution of other markets from the beginning..

COOK Report: So you don't have to have a master's or Ph.D. in Economics?

Franks: Well, I'm not so sure that a Ph.D. in economics really helps in an emerging market because many of the issues are social or political rather than purely economic. It does help to have a background in something technical, for example, an engineering background. I find it useful to be able to speak to the engineering concerns in an industry as well as to the concerns of traders. A background or knowledge in trading and risk management permits me the luxury being able to recognize what a product must be in commodity form and what industry infrastructural and technical issues may stand in the way of achieving this form.

COOK Report: Can you do what you're doing without needing to go into a fair amount of math with the people that you're working with? Do you have to staff up with specialist mathematicians and programmers who develop the algorithms?

Franks: It depends on the industry, but math and equations and theories essentially translate from one engineering discipline to the next. We sometimes give them different names, but the laws of physics are still the laws of physics.

COOK Report: So there's probably some mathematical or statistical tools that you use, but I'm getting the impression that what is needed is both the ability to work with technical engineers and experience in the processes of the development of markets that are common across ranges of commodities.

Franks: Process knowledge in market development is essential. It really helps having been there through the development of more than one commodity. Having been wrong about one particular market also helps. In its early days of the natural gas market, I had a bit of trouble initially understanding how it was going to work because of the transportation issues involved. I was one of those who said it couldn't happen. But the transport problems were solved and I had to eat my words.

COOK Report: Can you tell me a few words about that?

Franks: Sure. Natural gas took approximately 25 years to go through the commoditization process that began with the calls for industry deregulation in the late1960s. So it was a very long process.

COOK Report: Had oil already gone through the commoditization process?

Franks: Yes. There's been a crude oil futures contract since 1983. Natural gas didn't really start trading as a futures contract until '91. But there was a 25-year process that took us from a regulated environment, to one of a free and open marketplace. There's some positives to that long process. It allowed the industry to mature. In the beginning the industry was staffed principally with engineers, and engineers in that era were god. We did everything. We drilled the well, we built the pipelines, and then, when things started opening up, we did the marketing. . Eventually when it became clear that, as an industry, we needed to learn something about futures and cash market trading, we did that, too. Engineers were allowed and even expected to do pretty much anything and everything.

During this 25 years, the industry went through those changes that naturally occur when there's a disruptive event like the advent of price volatility. In the early 80s we witnessed the consolidation and shrinking of our experienced work force. The experienced work force was replaced with bright, young people, many of whom were fresh out of school. These young people really hadn't had the opportunity to learn the fundamentals of the industry such as how natural gas moved from point to point, or how it was bought and sold. They also had no preconceived ideas about barriers to change.

This deregulation, consolidation, and emergence of a commodity market occurred over a long enough time frame for the human beings in the industry to assimilate the change. During this period of time, the work force changed from one that worshipped the god of physics to one that worshipped the almighty dollar. Whenever there is an economic incentive for change, it will occur in spite of what may appear to be formidable barriers.

Basic Criteria for Commoditization

Now so much for background. Let's start out with the premise that, for something, anything, to be a commodity, it has to meet two basic criteria. One, it must be consumable. And the other very important issue in this commoditization process is that, whatever you create to trade as a commodity, must be fungible.

COOK Report: Easily used by a large number of people who would like to be able to use it. And uniform in quality?

Franks: Whatever it is that you're trading must be uniform and readily exchangeable. Like a natural gas molecule. You put one in the pipe, it doesn't matter if you take somebody else's out. And use it. They're the same; they're fungible; they're exchangeable. Now, that was a complication with crude oil, but the industry came up with a solution.

The complication arose from the fact that there's a great number of different grades in crude oil. Nevertheless it's all basically crude oil. What we were able to do, as an industry, is to create a way to equate one grade to the next grade. We took one grade, West Texas Intermediate, and began to trade that as a benchmark. The price of all other grades of crude was based on WTI with a pricing factor worked in.

COOK Report: Right, because you can still refine all of them into gasoline and presumably you get some various byproducts from one grade or various percentages that you don't get from another.

Franks: Or you may have to clean something out of one, like sulfur content. Crude oil was made to be fungible by designating one product as the benchmark or the reference point by which the prices of all other products would be calculated.

Ok, to get back to eating my words. Let's turn then to natural gas. Having gone through the experience with the oil industry,. I recognized that for natural gas to be a commodity, a tangible would need to be de-coupled from the transportation of that tangible - just like crude oil was decoupled from the transportation of crude oil to create a product to trade. With crude oil, it wasn't that hard to extract, because its transportation system is common carrier. With natural gas, the transportation was an essential component of the price prior to commoditization and varied from pipeline to pipeline rather dramatically.

COOK Report: In other words, what part of the cost of the commodity is the cost of the transportation to get it from point A to point B.

Franks: Correct. Correct. And you have to extract the tradable product from its transportation component . to have something that you can trade at a single location…. Like Cushing, Oklahoma in the crude oil market. The tank farm in Cushing is where West Texas intermediate is assumed to be when you're trading. When you go to natural gas, you've got a pipeline grid going just everywhere. And every segment is owned by somebody else who has different rules and regulations and different tariffs and a different way to price things. It's not common carrier. So I looked at that and didn't see how transportation could be decoupled since it was such a significant component of the bundled price. So, I actually said I didn't see how it would work. These are the words I had to eat.

COOK Report: So it's not common carrier in the sense that Williams has its pipelines and they don't intermix with Exxon's pipelines.

Franks: You have to buy the transportation separately from everyone and it's all structured differently. So it was a more complicated scenario to try to extract. Plus, the transportation of the natural gas molecule was a more significant part of the price than it was in relationship to crude oil. Still I had to eat my words, because those are stupid reasons to say it can't happen - but I was younger then.

COOK Report: And because obviously people figured out how to make it happen.

Franks: Well, natural gas is really fungible. There are far fewer grades of natural gas than there is in crude oil. That part of the equation was actually much simpler.

COOK Report: And it helped keep driving it forward.

Franks: And I had to eat those words about extracting the transportation cost from the tangible. The natural gas industry began to trade around a location in Erath, Louisiana, Henry Hub with a standardized product, under standard terms and conditions and standard quantity, and so on. It became a very easy thing to trade. The natural gas futures contract actually is one of the most viable contracts traded..

From Energy Products to Bandwidth

COOK Report: Well, how did you get from these energy products into bandwidth.

Franks: Me, personally?

COOK Report: Yes.

Franks: Okay, I was happily sitting around as a principal of Sabine, the owner and operator of Henry Hub, having a lovely time at my job, when utilities in the electricity industry began calling. Some in search of information about hub services, while others wanted to recruit me. Henry Hub is premier gas market center. This was the early 90's. I was convinced by Portland General Electric to come to Portland to help create an environment favorable to the creation of a commodity market around one of their assets, the California-Oregon border. They hoped to attract NYMEX to that location and hoped it would become the delivery location for the emerging electricity futures contract.

COOK Report: And that must have taken you into Stan Hanks' physical area.

Franks: Into Portland, Oregon. Yes, it did.

COOK Report: And Enron was involved with the Portland company?

Franks: Enron was not involved when I began to work with Portland General, but they ended up buying Portland General a couple of years later. While I was working on the electricity market at Portland General, I was looking toward the future and watching for the next product that had the potential to emerge as a commodity. This emerging market focus is my little niche, to be there on the front edge of an emerging commodity. I had been looking at bandwidth as well as a couple of other things. In the context of bandwidth, Portland General had this lovely little asset that, at that point in time, was called First Point Communication. This is the asset that is now known as Enron Broadband Services. This looked like a good place to begin to add to my fundamental industry knowledge and to begin to flesh out the concept.

COOK Report: What exactly did it do?

Franks: it was a start-up that was beginning to build a communication network. It is now called Enron Broadband Services,but back then it was called First Point Communication and had a different focus. Earlier that year, Stan joined the communications side of Enron and changed their world, technologically speaking. I was ready to move out of the electricity market and talked with Joe Hirko, who has since become the CEO of Enron Broadband Services, about researching bandwidth commoditization. As it turned out he was in favor of looking into it.

COOK Report: And the idea of First Point was that the raw material was a fiber network of some sort or another? And they wanted to figure out how to use it and develop it?

Franks:At that time, they were planning to build a communications network, but it had a very different focus, a very different architecture than what would result from Stan coming on board. Joe wanted to elevate the perception of this asset with the company. Since Enron is a market maker, they would naturally be interested in a new market that had the potential of leveraging assets they already owned. So, Joe agreed to start looking to see if it made sense. As a result, I did the initial internal work to write up what could occur and what some of the economic possibilities were.

Hanks and Franks Join Forces

Joe introduced Stan and me in early 1998. And within 30 seconds, we were really good friends. Stan had begun to look at the

commoditization of bandwidth about a year previously, as had I. So, independently, we have both looked at it for about three years, together for two years. Our knowledge and experience is very complimentary. And while I have no way to qualify this, I really believe that means that we've been looking at it longer than probably anyone else in the world.

COOK Report: What did you and Stan find, what were some of the next steps?

Franks: Well, I'll tell you, the first step, was for Stan and I to kind of lock ourselves in a room and exchange information. I had no doubt that he was (and still is) the most technologically competent human being in my personal network of friends and advisors. He, in due course, made sure I had all of the details that I needed to have and filled in the areas that were missing in my Internet technology knowledge database. I filled in his knowledge base with supplemental market development and commoditization information.. What I needed to do was to explore the physical aspect and the infrastructural aspect of this marketplace in order to figure out where a good starting point was and to formulate how we could make something fungible to trade. It was necessary to look for barriers and to conceive of a way to handle them.

I also did research around the history and the regulatory environment worldwide. Basically I had to make myself a book of everything significant about this marketplace…. how bandwidth is bought and sold, what the competitive elements are. When I start through this process, I'm looking for a tangible to extract from the transportation of the tangible. But it can't be done exactly that way with bandwidth, because what there is to buy and sell is the transportation.

So, not only did I want to find something that I could extract, which I really couldn't, I also wanted to find something that was technologically independent. There are many different philosophies about which is the right way to go technologically speaking. ATM? IP? Both have significant attributes. It's a big, emotional debate.

And if whatever commodity developed had a characteristic that made it only feasible on one of those technologies, then it wouldn't work. Because you've got to have something that's technologically independent, that does not favor one particular industry participant over another. Consequently we're looking for something that everybody can embrace and doesn't look like it has the flavor of one entity over another.

Stan and I spent quite some time thinking through various scenarios. Stan clearly has a very good business mind and clearly understands the financial side of the world, which is highly unusual in a person with his technology background. Nevertheless, we put a few communications peoplethrough a little mini Futures 101 class. I have found that if people get some of that kind of experience, it makes it easier for them to come up with ideas and notions about things that either will work or will become stumbling blocks. It helps move the entire the process forward.

The exercise we went through was really very low tech and very basic . I gave the different participants in the room a scenario. Each had a role to play. One worked for XYZ Company and as responsible for hedging his company's assets. Some people were speculators, some people were hedgers. I gave them a real scenario about the level of comprehension that their own board of directors had about what they were going through, which was hedging. And so they had to deal with those real issues while trading.

COOK Report: What kind of trading?

Franks: I think I used natural gas. I varied the fundamental events and issues while they were trading, because this is the kind of environment that someone running a risk management desk, for instance, has to deal with.

COOK Report: So you were talking in terms of the processes from the earlier commodities markets? You were not trying to deal with TCP/IP bandwidth?

Franks: Right. All the people in the room had some knowledge of the technology space in which we were trying to work, but what they lacked was a pretty good comprehension of what a trader has to go through. So they needed that piece of information so that it would stimulate their creativity in coming up with possible scenarios for things they could trade in bandwidth and what some of the barriers might be. You do all this because you don't go way down the development path only to discover that physically something can't occur around which you had built your entire plan.

COOK Report: Sure. What were the kinds of things everyone learned?

Franks: Well, they got an introduction into the type of mind that a trader has. Traders are not into detail. They are not into long explanations and calculations about things. They need something that can be traded just based on price differentials. So that they don't have to think, they can react to the market.

For fun, at the beginning of the exercise I predicted, knowing a little bit about the personalities in the room, who would actually complete the task I had given them and who would not. And the person I said would actually get all of his trades completed was the person with the least knowledge about technology or natural gas. It was the person who was fresh out of getting an MBA, and who could just do what was told to him, without thinking. People like Stan are going to spend too much time analyzing. They're not going to finish.

COOK Report: Where would these hypothetical traders work in a bandwidth trading world?

Franks: There will be a couple of philosophies. And everybody will begin to debate this, by the way. They'll debate is it better to bring people in from the risk management community with, for example, a background in energy risk management and teach them the fundamentals of the industry and how to trade it? Or is it better to take people who actually are buying and selling bandwidth as a normal course of business in the traditional manner, and train them in energy risk management and trading.

COOK Report: Do you have some opinions about what the pluses and minuses are, from your point of view?

Franks: Of course I do, but then you have to realize that, in fact, this market does appear to be developing at a very fast rate of speed — and you remember I told you natural gas took 25 years? The market that came after that one took only three; that was electricity. And if that trend is followed with bandwidth and I have every reason to believe that it will be, that means that this market will be quite active and viable by the end of this year.

Training for an Active Market by Year's End.

COOK Report: Help me understand how we get from here to there.

Franks: Those who are beginning to be focused on the creation of the staff to do this, they can debate those issues all they want. But for everyone else the reality is, where are they going to get the resource that's accurately trained or adequately trained in how to do it?

COOK Report: If a big carrier is going to buy and sell bandwidth in a commodities market, somebody in that carrier has to be trained in what way? To do what?

Franks: There you go. That's why we have put together a comprehensive training program to meet this need. You've got to be there when your customers believe they need it and these things do take time to develop. Companies, even those who are on the edge

of this market like Enron, as talented as they are in risk management trading, and as obviously competent in the area of communications as they are, are still faced with the same issues in staffing that others are. In this situation, they may bring people in from their own energy risk management trading organization and, indeed, have done so. They're also going to hire or bring in people with the fundamental knowledge of how bandwidth is actually bought and sold today, traditionally. But they're also going to hire bright young minds, bright young MBA's.

In this sense Enron is no different. This is what everyone is going to have to do, they're going to have to mix it up. When you bring in a group like that and you're hiring and staffing and need to train with the perception that this market is going to be quite active by the end of the year, then you've got to lean on people like those at Andersen consulting to assist you in your training program.

COOK Report: So what exactly are we training for? I had one image of where you might have the development of a bandwidth commodities market that would be like a stock exchange. In other words something that would be an independent organization in itself. Such an organization would have its own employees. People would connect their computer terminals into it and they could see prices. Everybody would plug in and if I wanted to check a DS3 availability from January 1 of next year through the end of that month from point A to point B, I could plug into that market, look at the database, see what the prices seemed to be to get any ideas about whether there were concerns about availability. And I could say, well, I want it and I'll pay such and such price.

Franks: Let me stop you right there. That may be something that occurs in this industry. When I say this market's going to be viable to trade by the end of the year, that does not suppose that there is an electronic platform that everyone uses to do that by year's end.

COOK Report: Because it sounds like the initial stage is going to be at each major provider that you're going to have one or two or three or X number of employees of that provider trained in the skills of how to talk with a number of similarly trained employees at some other provider. In that sense, it's going to start out all very decentralized.

Franks: Well, and I'm not going to even predict whether it becomes centralized or not. Because that isn't the most material issue. This market is at the very beginning of development. It is highly unusual to see it reach this kind of fever of interest when the actual benchmark product isn't yet determined. Furthermore, every one of these earlier markets has developed as a result of some industry disruptive event, except for this one. Except for this market, all the other markets created something to trade around, once there was price volatility. This market doesn't yet have what we would call price volatility. The price is not going up and down, it's just going down.

Supply Unlikely to Keep Pace with Demand

The price is in an exponential decline. And quite a steep one. But the real disruptive event is the realization by those in the industry that, regardless of technological advances, regardless of their fever to lay fiber across the world, there is really no way that supply is going to consistently be able to keep up with the ungodly increase in demand at which we are looking. Whereas there may appear to be a supply glut right now, bandwidth demand will rise to fill it and will exceed the available supply. Then there will be another technological advance that will increase the available bandwidth supply. Then demand will rise again and so on and so forth.

COOK Report: But with the technology advances, what you can squeeze out of a piece of fiber is accelerating.

Franks: But when you look at the fact that the Internet is becoming very rapidly the principal communication tool of the world market and particularly world commerce, it's just a given that you have to have an e-commerce site if you are a business. Applications are created on a daily basis that require more and more bandwidth and a population that expects instant gratification. Large quantities of data downloaded in a blink of an eye. Software run over the Internet with a speed previously enjoyed only when run on desktops. You have a population that has grown up on STAR WARS and STAR TREK and has very little true comprehension of what is do-able and economically feasible. And their expectations are without boundaries. Consequently, demand grows to fill available supply..

COOK Report: So you have a need for risk management and financial planning in almost every level of every player in the industry, yes?

Franks: There are a great many risks that need to be managed in this emerging market. For example there is significant capital at risk when one elects to build a circuit. As an indsutry, we've got to do something infrastructurally to make it possible to buy and sell excess bandwidth for shorter periods of time to better utilize the assets already in place. The creation and development of this incremental market will provide a fway to do just that.

COOK Report: And again, probably much more so than with any other possible commodity, the rate at which you use TCP/IP is difficult to predict because it is so very bursty.

Franks: Absolutely.

COOK Report: All of a sudden, I need everything that's there and then all of a sudden, I go relatively quiet.

Franks: That's right. So what that means is that one thing this market, whatever shape it takes, will provide relatively quickly, is the ability to buy and sell some standardized form of bandwidth for shorter periods of time without the weeks to months it take today to provision it. This is a very positive thing. It will give traditional people in the industry an additional market and additional people to buy and sell what they are currently unable to.

COOK Report: It sounds like there is an educational mission or need that you've just outlined. You want a cadre of people who understand the implications of all this. Given the players out there presumably it's going to start with the identification and training of some key people in most of the key dark fiber providers. These people will then be able to interact with the key people that have been similarly trained in the clients or potential clients of the fiber providers.

Franks: It usually starts at a fairly high level inside the organization. It involves letting management and its executives truly comprehend what this market is and what it may or may not be capable of doing for them. Because many of the risk management-trading-broker types that are coming in and speaking feverishly about this issue sound to the traditional telcos as if they are oversimplifying the situation. It also sounds like they're trying to change all the infrastructure of the industry and change all the way business is done and that isn't actually the case. Well, we have two different cultures trying to speak two different versions of English. So there is a level setting that needs to be done. Everyone needs to understand what it is that they are actually saying and what the market can become.

COOK Report: Then the big problem for the carriers is how they will make sure they have enough capacity? But on the other hand, not too much capacity?

Franks: Yeah, everybody must have their own philosophy on this. But this is one area where everyone can come together and are doing so. Still it's also confusing to many of these people when they see reports like the

one Schroder & Company came out with giving a very positive valuation to Enron, just because of the existence of bandwidth trading. And while there was nothing negative about what the report said, if you're not well aware of where this product is in that stage of commoditization, you can draw incorrect conclusions.

COOK Report: So presumably with every company or person that you want to talk to and interact with there's no difficulty in everybody recognizing that we've got a problem here. You can come in and say, "I want to talk to you about the problem of supply and demand and pricing and the risk management policies, that that you have. I want to talk to you about some ideas as to what you can do to plan more rationally." And at some point, you can easily agree with each other on what the problem is.

Franks: That isn't an issue. People are very hungry for information at this point. And also for ideas and knowledge.

COOK Report: Then you come in as someone who says to them: for the following reasons, it looks like the industry needs to move in this direction and here are some ideas that will help you understand why and what needs to be done.

Franks: Well, I don't usually take a position as to whether the industry actually needs to move in this direction, I'm just there to say here's what it is, here's where it is, here's what's going on, here are some possibilities. Here's what will or will not resolve. I do that at a fairly high level and then if they decide they want to go further, we can. There are a number of different roles one can play in this market. And many times companies need an assist in strategizing their entry into a market and, specifically, which role they should be playing.

COOK Report: It sounds like in the one sense that there is probably a segment of companies out there where the educational role is very basic and just beginning. Probably some other portion of the companies presumably you've already been talking to and various sets of people are getting more aware. But what do you have to do the rest of this year to get from this initial everybody knows we've got a problem, to well, since there is a problem, some of the ways of looking at the problem are to think about how would you develop bandwidth trading. If I'm an executive and I listen to your presentation, I'm sold, what happens? How do we get from the very beginning to where you see things being at the end of the year/

Franks: Well, each company is going to have to make a determination as to what they want to see at the end of the year. Which we can, of course, help them figure out, based on their core competencies and risk profiles and basic philosophies. But then it's very important and we do advise clients to get involved in industry endeavors to create, for example, a bandwidth trading organization, which will be a self-governance organization, for issues surrounding bandwidth trading. We encourage them to participate, if they are not already, in any kind of industry endeavor or task force around Quality of Service. And any of those issues that fit into what role they believe they should play and what they would like to see in their industry. They just have to be active.

Straw men in Search of the Industry Benchmark

COOK Report: If you're going to have some kind of a market functioning, you have to develop, between now and then, you have to develop some kind of uniform contract, you have to develop what your benchmark is and what's going on there? Is there already a nascent industry trading organization beginning to form?

Franks: Yes first meeting is actually next week (March 23). These things are happening rather rapidfire, but the first meeting of that is actually next week, in Washington, D.C., and there are carriers and representatives from energy trading groups. For the Washington meeting there is already one straw man out there that Enron has proposed, but there are other straw men under development that will be thrown around in the industry as possible benchmark product to trade. Now about three weeks ago Stan talked to you a bit in your interview with him about an OC48 product.

COOK Report: Yes.

Franks: In the meantime he and I have discussed it further. It makes sense to you talk about an OC48, but it's harder to see how to use a benchmark product if you limit that benchmark or design that benchmark around a specific city pair. Enron's straw man, or the initial thing they're throwing out there to trade, (not that they believe it will be the ultimate benchmark,) is a DS3 between New York and L.A. If one continues down that path and stays with a specific city-pair as a benchmark, I fear that it makes the product, or the benchmark, more difficult to apply to other city-pairs. Like London to Tokyo.

COOK Report: Okay. Are there some alternative ways of thinking about the situation.

Franks: What Stan and I came to this week is, yes, perhaps an OC48 makes a great deal of sense, given the way technology is moving. Maybe not today, but at some point in the near future. Again, note that, we're trying to be technologically independent. But We believe that should be a mileage-based benchmark. For example,an OC48 contract-size 100 miles. Now, that's one step. I f you take it a little further down our evolutionary path, and perhaps beyond this year, because the industry is going to have to learn how to use this product, then you pull the electronics off of the OC48 and what you end up with is a lightwave.

COOK Report: Like a color in the sense of dense wave division multiplexing.

Franks: Right, so it is quite conceivable in the near future that what we will be trading is some quantity in miles, maybe 100 miles is too small, it might need to be larger just to have a decent-sized contract. But perhaps we'll just be trading lightwaves. Which also has a cool factor. Who wouldn't enjoy having a business card that said "Lightwave Trader".

COOK Report: In the sense of advanced technology?

Franks: Yeah, so I mean, your business card, instead of just saying Futures Trader or Trader, says Lightwave Trader. Trader of Lightwaves. That's somewhat futuristic But it actually works. It then becomes that thing that Stan and I were trying to do two years ago in coming up with something that is technologically independent and is not tied to a city-pair. It becomes the fungible product we're looking for.

COOK Report: Just to underline something then, at the risk of being a little didactic, if you're trading an OC48, that is, in order to be an OC48, you have to have specific equipment attached there. And you're already slicing and dicing a lightwave in certain ways. So what you are doing is moving a little bit up in the level of abstractness.

Franks: Right, but it doesn't make sense to begin to trade lightwaves right now, because it doesn't solve anything for the industry. These developments need to occur as the industry can absorb and accept them….and must be the result of a need in the industry to work.

COOK Report: That's true. But suppose you're thinking of a short-term, immediate need of two weeks from now and there's a convention in such-and-such a city or some event and we need to get into point A and out of point A with some sort of quantity.

Franks: Well, lightwave trading isn't what you can do right now. If it happens at all it is some point in the future — and it may be just two years away, given how fast these things tend to evolve, then it does make sense to trade light waves, because we have learned much, much more at that point in

time.

To take it back to your concept of an electronic trading platform. Keep in mind that the movement to an electronic platform does not change business, it only speeds it up. Also keep in mind that this bilateral trading market has not yet developed any kind of consistency that makes it actually possible for anyone to put what they need to put on an electronic platform. This situation makes it very difficult to conceive of a single electronic platform for the buying and selling of this form of bandwidth market. It is more likely that multiple electronic platforms for the buying and sellingof bandwidth will emerge….some may be for bilateral markets and others for blind match…others for auction and reverse auction. It is also easy to visualize a trades subscribing to multiple platforms….buying one then selling another….continuously taking advantage of the opportunities.

I do believe that after this year, we will see a lot more reason to have perhaps maybe two or three different electronic platforms that become, then, portals for the carriers' storefronts. And it makes it a seamless way for industry participants to go out and buy and sell bandwidth. But, we also need to have some consistency between products, we need to have some standardization of terms and conditions. And some actual reason to do it on an electronic platform. If it otherwise takes you 90 days to get started after a trade, then the speed of an electronic platform is hardly relevant.

COOK Report: Well, technology developments are coming. For example some of the stuff that Sycamore is doing, with its new optical switching.

Franks: Absolutely. But they're not there yet.

COOK Report: But, six months, a year?

Franks: Absolutely. However, we've got to have those developments to make it even necessary to use an electronic platform, otherwise we're force-fitting the electronic platform on something that works just fine on the telephone.

COOK Report: One question. Is the idea that, if you look in the most general way at what may happen between now and the end of the year the educational process that you're undertaking with executives in various and sundry companies will continue. You're seeding at the periphery and as a result of the seeding you're doing, presumably companies are going to begin to get interested in joining this incipient bandwidth trading industry forum?

Franks: Well, I don't think that it's just education through the end of the year. Enron has already thrown the straw man out there for the DS3 and completed a few transactions under that. These are transactions with other industry participants. So, in addition to getting educated, I think we will see people jumping in and getting their toes wet.

COOK Report: In other words transactions may just be something that individual participants in the forum decide to do on an individually negotiated basis. Let's experiment over the next months. I'm carrier A, you're carrier B. I would suggest that when we have a need or something to trade, here's some defined basis on which we could consider contacting each other. Here's what we could do and here's maybe how we could do it.

Franks: Right.

Use of IGBP in Network Design Disscussed

Many Variables and Tradeoffs illustrate Complexity of Design Decisions

Editor's Note: We publish a narrative summary of roughly the first half of a discussion on the use of IBGP protocol in establishing a network mesh architecture with more that 30 Cisco 7500 series routers. Among the design issues involved are the ability to deal with fast convergence versus increasing changes in network "state".

On April 1 on the inet-access list **Andy Walden** wrote:

Just curious about some of the issues that could arise through putting 35 or so Cisco 75xx at pops on a major providers backbone and then running IBGP with a single AS and BGP with the provider and probably some others also at each point. What is a safe number of peers for fully meshing like this? Would confederations or route reflectors work better to cut down on the size? How many routers in a confederation would work best? What else needs to be considered in a situation like this? Thanks for the time.

Dave Cooper: About35 routers w/ 75xx should be no problem running a full ibgp mesh. I would only consider route-reflectors if you plan on expanding further down the road. If your running client-to-client reflection, this will limit your client's number of *best paths* to the number of route-reflectors. This and other factors need to be considered before running rr's.

Vijay Gill (April 2): Of course, it could be argued that BGP is for disseminating NLRI, while the path selection (which may or may not include ECMP), should be handled by the IGP. Since the peers only propagate best paths and the clients are downstream of the RRs, one could argue that the multiple view at the client level isn't buying you a whole lot.

Sean Doran: This would be a silly argument.

All routing protocols disseminate NLRI; that is their function. This disseminated NLRI is then processed by path-selection algorithms in various individual routers, in the process of building a forwarding table.

Any routing protocol presents either an indication of reachability of prefixes through one's immediate neighbour, or data with which one can construct a connectivity graph. The former is a vector, which is a bit of information along the lines of someone telling you, "I know how to reach this network, and it costs X to use me". The latter is a link state, which is like, "This link is working, and using it costs X". Link-state announcements also usually carry some "undefined-cost vectors", which behave like, "These prefixes are directly reachable via this node". (The prefixes may be directly connected, or reachable over some "foreign" network the IGP cannot map).

NLRI = Network Layer Reachability Info, which is one of those wonderfully easy to understand BGP terms that basically means "vector". So, you're partly right — iBGP in a sense interacts with an IGP to provide the equivalent of undefined-cost vectors. ("These networks are directly connected to node X", where X is usually a loopback or interface address).

So, I don't think you mean "path selection" here, since that means sorting through a RIB (== Routing Information [Data]Base) filled with information from a particular routing protocol. IGPs and EGPs (like BGP) simply populate a RIB. A path selection algorithm is run by each router on that RIB, in order to determine the appropriate to program one's forwarding tables.

For routes in the iBGP RIB, the path selection algorithm results in a BGP next-hop address, which may not be local. One therefore does a path selection using the IGP RIB to determine how to reach that next-hop.

Path selection algorithms are also used to construct BGP adj-out RIBs, which are essentially lists of vectors one will send on to one's BGP neighbours. The IGP RIB is usually unimportant except occasionally as a resource for extra information (e.g. outgoing MEDs).

I don't even know what ECMP *means.* In English, maybe?

Gill: Equal cost multipath. [. . .] pathing was the incorrect word to use.

As far as ibgp meshes go, 75 routers in a full mesh is trivial with the caveat that routers Dave Cooper is running tend to be on the high end of the scale as far as horsepower goes.

Cooper: I would not suggest running confederations.

Gill: Reasons?

Cooper: As far as other things to consider... RSP-?, memory, number of full views each router is taking (which will bog memory and cpu when converging).

Gill: When doing resource consumption modeling, it has turned out that CPU isn't quite a problem, but rather, the propagation of the tables is. That is, quadrupling your CPU does not help 4x. At some point, BGP, being carried by tcp with attendant slow-start, congestion control etc. mechanisms, runs into a convergence issue as your updates asymptotically approach 100 as a percentage of the time on a percentage basis as your cpu power approaches infinity.

Doran: "a percentage of the time on a percentage basis"? Are you a civil servant?

Gill: Just coming back from Sweden, speaking in tongues.

Doran: I _think_ this sentence fragment reduces to:

As your CPU power increases, the amount of CPU time spent doing updates gets closer to 100% but never quite reaches it.

Wow. How about rephrasing this into something that makes sense?

Gill: In simple words: your upper bound on convergence is not cpu, but rather how fast you can spread the knowledge around. All routing protocols currently work on old state, the limiting factor is getting that state spread around

Doran: Well, if you can find a routing protocol that can predict future state reliably, rather than propagate stuff it already knows about, I am sure you will find much more interesting applications for your algorithms than simply doing IP routing. -:)

If you are referring specifically to BGP-over-TCP, rather than any other protocol, then this is likely the case. The limiting factor here is TCP, so in the presence of unlimited CPU and memory bandwidth, we are stuck with BW = max(cwnd,rwnd)/RTT.

However, there are all sorts of assumptions in there; in reality, you will never eliminate all bottlenecks in a sufficiently complicated network, and while in steady state you may have sufficient resources to make max(cwnd,rwnd) interesting, it is not at all clear that CPU and memory bandwidth will not be the bottleneck in the event of Byzantine de+reconvergences. This is, after all, one of the reasons TCP is used for BGP - slamming rwnd down to zero is a perfectly acceptable response by a resource-short receiver.

Moreover, the world is not BGP; a non-TCP optimal flooding scheme should be able to propagate all state changes at propagation delay (you insert the state change information at the front of any outbound interface queue). By giving control traffic precedence, and by being clever in queue maniupulation (rewriting already-queued state announcements), you certainly can approximately eliminate the argument that spreading state knowledge around is the limiting factor. However, now you have to begin to worry about the amount of CPU time spent in processing all the state knowledge, and what to do in the event you have a resource-short receiver...

Fundamentally, fast convergence is an attractive deity for some to worship. I don't pray at that altar, because I believe I know how to scale a network for stability, rather than for increasing amounts of state change. That is, in my opinion, resources are better spent on keeping network state change to a bare minimum, rather than dealing with increasing state changes over time. Of course, this is a pragmatic engineering decision; the correct answer is a good research topic. -:)

Gill: If you take time X to reduce a set of nlri into a forwarding table you can use to determine what out interface to send packets to and to spread that knowledge around so you have "convergence", the faster your cpu gets, the time X is bounded by dissemination of information. i.e. if you get infinitely fast CPU, your convergence will not approach 0 s.

Doran: Well, yeah, because of propagation delay. But you mean "will not reach", right? (Hence, the original use of the word asymptotically, yes?)

It certainly can approach 0s, simply by making a network occupy a very small volume (thus reducing propagation delay).

Effectively, what you seem to mean to say is that if you throw enough resources at the convergence problem, you can make it work such that even huge network instability can be managed, if there is sufficient inter-router bandwidth available for control traffic. Or am I misunderstanding you again?

Gill: I am saying quite the opposite, that throwing resources at a problem is NOT going to work (resources in this context mean super fast CPUs), hence need for aggregation and hierarchy.

Dan Golding: There was certainly a time when 75xx routers were considered to be on the high end of the horsepower scale….

Doran: Only by Cisco marketing, and their apologists.

The 7500 was initially a DOWNGRADE from the 7000 + SSP. Rather than continuing the intelligent separation between the packet-forwarding function and the router function (which programs the packet-forwarding hardware), everything was glommed back onto a single CPU. It was not until some time later that the Wellfleet model of intelligent interface cards, which did their own forwarding independently from the router processor (but under its direction) that the 7500 became anything but an AGS-.

There are still horrendous flaws in the 7500 architecture. The busses are really badly thrown together, and the MEMD mechanism for absorbing short-term bursts towards the "switching fabric equivalent" is insufficient to support high bandwidth traffic at any reasonable end-to-end delay.

A much smarter thing to do at the time would have been to build a bigger, better SSP, and evolve towards a multiple SSP system, where there would be some sort of SSP-to-SSP fabric, with each SSP driving a set of stupid interface cards.

The VIP model is much the same — you get an SSP-like function in the VIP side, and they drive one or more PAs.

The GSR model is a cleanup of the VIP model, or rather a complete redesign. It works reasonably well. What a pity that you need 7500s to talk to a huge number of interfaces, because Cisco's focus on line cards differs substantially from what the port adaptor people look at.

Golding: But that is certainly not the case now. They are mid-level at best. The big peering routers these days are Juniper M20/40, GSR boxes, Foundry BI8000, etc. It's awfully easy to peg the CPU on a 7513 with RSP4.

Doran: What are you doing with the RSP? If you use it to forward traffic — i.e., you're not doing distributed CEF, and/or you're using old interface processors — yes, I'm not surprised the CPU pegs. But otherwise, the CPU should be sufficient.

Golding: I would be cautious about running 75 routers in a full mesh when there are other solutions.

Doran: Why? There are tradeoffs one way or the other; a full mesh exposes fewer hard-to-fight operational conditions (and fewer bugs!) than most of the other solutions

Some ICANN Footnotes:

Editor's Note: ICANN's story doesn't change. It just continues the charade described in the last issue - Internet governance by Marx Brother's comedy. Some of the amusing skits that we have noted since the May issue follow. The last of these four items deserves a bit more respect. It is John Berryhill's rather brilliant commentary on an amusing skit.

Footnote 1: NEW INVESTIGATION

Last October 19th buried in a Conference report in the closing days of congress and unnoticed by the press, another ICANN investigation was ordered, this time by the GAO.

CONFERENCE REPORT ON H.R. 2670, DEPARTMENTS OF COMMERCE, JUSTICE, AND STATE, THE JUDICIARY, AND RELATED AGENCIES APPROPRIATIONS ACT, 2000 (House of Representatives - October 19, 1999) is the title of the report. http://www.osec.doc.gov/bmi/budget/Legis/FY_2000/CJSBill,Report.html is the url.

"The conferees direct the General Accounting Office to review the relationship between the Department of Commerce and the Internet Corporation for Assigned Names and Numbers (ICANN) and to issue a report no later than June, 2000. The conferees request that GAO review: (1) the legal basis for the selection of U.S. representatives to ICANN's interim board and for the expenditure of funds by the Department for the costs of U.S. representation and participation in ICANN's proceedings; (2) whether U.S. participation in ICANN proceedings is consistent with U.S. law, including the Administrative Procedures Act; (3) a legal analysis of the Department of Commerce's opinion that OMB Circular A-25 provides ICANN, as a `project partner' with the Department of Commerce, authority to impose fees on Internet users for ICANN's operating costs; and (4) whether the Department has the legal authority to transfer control of the authoritative root server to ICANN. In addition, the conferees seek GAO's evaluation and recommendations regarding placing responsibility for U.S. participation in ICANN under the National Institute of Standards and Technology rather than NTIA, and request that GAO review the adequacy of security arrangements under existing Departmental cooperative agreements." Editor's Note: Apart from its contribution to the general American strategy of obfuscation where control of DNS and hence the Internet lies, we don't have great expectations from the study.

Footnote 2: THE EUROPEAN COMMISSION TRUSTS U.S. LESS

At PERLINK http://www.icann-studienkreis.net/intgov.htm http://www.icann-studienkreis.net/intgov.htm we find dated Brussels April 7th the following report. Bold faced emphasis is our own.

COMMUNICATION FROM THE COMMISSION [of the European Communities] TO THE COUNCIL AND THE EUROPEAN PARLIAMENT

The Organisation and Management of the Internet -- International and European Policy Issues 1998 - 2000

EXECUTIVE SUMMARY
Introduction

The organisation and management of the Internet infrastructure involves several limited but essential technical coordination functions. This Communication addresses: - recent developments in this area, during the 1998-2000 period. - transfering the US Government's responsibilities to ICANN* - the principal policy issues for the European Union and internationally, and - operational conclusions for the European Union. The Communication also draws attention to the current expansion of the Internet in Europe and its increasing importance as a key economic and social infrastructure. This is likely to put the capacity of the existing system under some strain. The Commission has already taken some measures aimed at improving the economy and efficiency of the communications infrastructure for Internet use and will continue to monitor the development of the situation.

International Aspects

The European private sector participants have played a critical role in establishing the European Union's position at all levels in the global coordination of the Internet infrastructure functions: the ICANN Board and Supporting Organisations, the DNS Root Server system, Internet Registries and Registrars and in the IETF and the World Wide Web Consortium. Without that commitment, the public policy role of the EU and the Member States would be much less effective, if not impossible. Maintaining and deepening European private sector membership and participation in the ICANN organisation is a critical pre-condition for successful participation by the EU both from the point of view of the Internet user community in Europe and from the point of view of public policy.

Regarding the EU's international role, the Commission requests the Council and the European Parliament to confirm the Union's existing role as a participant, co-ordinator and, where necessary, negotiator in this area. This involves the International Organisations, notably WIPO and the ITU, bilateral relationships with several governments, including the United States and the role of the European Union and the Member States in the ICANN Governmental Advisory Committee (GAC).

These international responsibilities go beyond a simple presence and oversight of the ICANN process. They also involve specific aspects of EU public policy: - the neutral global role of ICANN - the scope of the US Government's remaining powers over the Internet infrastructure - international aspects of intellectual property, competition and data protection policy - the scope of ICANN's authority regarding Registries and Registrars. **Creating and maintaining an environment for neutral international jurisdiction is proving to be even more difficult than had been originally envisaged when this process began in 1998. [**Snip]

Domestic EU considerations

The Communication also addresses a number of European policy issues that require further attention during the months to come, that will also require the continued support and cooperation of the Council and the Member States if the Commission's objectives are to be achieved. These include: – Internet Domain Name System: appropriate follow-up to the Commission's recently launched public consultation on the creation of a new Internet Top Level Domain: Dot-EU. An additional Communication is envisaged on this question before July 2000; – Intellectual Property Rights: preparation of a code of conduct or other appropriate instrument to address abusive registration of domain names("cybersquatting"); – National Country Code Top Level Domains (ccTLDs): - implementation of the guidelines recommended by the ICANN-GAC - preparation and implementation of guidelines for data protection and privacy - development, by the national Registries of best practice for registration policies. – Alternative dispute resolution (ADR): Development of and implementation ADR policies in the light of the WIPO recommendations appropriate for TLD Registries operating in the EU. –

Competition policy:

The Commission will ascertain whether agreements and business registration practices in the area of Internet Organisation and Management fall under the EU competition

rules (Articles 81 and 82) and, where necessary, will take the appropriate action on the basis of its direct powers under the EC Treaty". These policy developments will continue to be co-ordinated as appropriate with the Member States through the existing Internet informal working group, convened by the Commission.

The Internet Infrastructure

The topography and capacity of the Internet backbone infrastructure in Europe is a source of some concern. **The current structure of prices and available bandwidth have had the effect of diverting a significant proportion of European Internet traffic across the Atlantic and back. The resulting costs and inefficiencies are already burdensome and will become intolerable as increasing proportions of communications and commerce migrate to the Internet in the foreseeable future. The security and competitiveness of Internet communications in Europe consequently depends on the security and costs of the US-based Internet exchange points.** The Commission intends as a matter of urgency to complete its information in this respect and to pursue its policy aiming at encouraging the rapid roll-out of very high bandwidth Internet backbone throughout the European area.

Conclusion

In the light of the Presidency Conclusions of the Lisbon European Council, the significance of these issues can only be re-emphasised. **Indeed, the whole scope of the Information Society and electronic commerce in the European Union, and world wide, is influenced by the stability, and reliability of the Internet in the context of its extremely rapid growth.** The European Parliament and the Council are invited to endorse the policies and actions envisaged by the Commission in this Communication and to support their implementation in co-operation with the Member States. These are summarised in Chapter 10 of this Communication.

[And in section 5.4 report itself we find the following chain of reasoning. (a) DoC has reasserted its control over ICANN. (b) the European Commission has been assured by DoC that ICANN will receive a full turn over of US gov't authority by October 1, 2000 [**Editor's Note**: if the commission really thinks the Clinton Gore Administration will divest the root less than 40 days before the US elections, it is seriously deluded.] (c) the commission expects necessary governmental oversight of ICANN to be achieved on a multilateral basis, in the first instance through the Government Advisory Committee.]

5.4 Exercise of ICANN's Authority

"Under the agreements with US DoC and NSI, ICANN is also obliged, to the Registry and to all accredited registrars, to comply with specified procedural requirements governing the exercise of its authority, particularly regarding consensus-building. The US DoC has also reasserted its rights of supervision over ICANN policies, including any amendments to ICANN's agreements with NSI. Furthermore, ICANN shall not enter into any agreement with any successor registry to NSI for the .COM, .NET. and .ORG TLDs without prior approval by US DoC. Should US DoC withdraw its recognition of ICANN or any successor entity by terminating their Memorandum of Understanding, ICANN agrees that it will assign to DoC any rights that ICANN has in all existing contracts with registries and registrars."

"The broad scope of the powers and authorities reasserted by the US Administration (as recently as November 1999) notwithstanding, the US Department of Commerce has repeatedly reassured the Commission that it is still their intention to withdraw from the control of these Internet infrastructure functions and complete the transfer to ICANN by October 2000."

"The Commission has confirmed to the US authorities that these remaining powers retained by the United States DoC regarding ICANN should be effectively divested, as foreseen in the US White Paper. The necessary governmental oversight of ICANN should be exercised on a multilateral basis, in the first instance through the Governmental Advisory Committee. The Commission will take the necessary steps to ensure that the principles of openness, transparency and respect for international agreements, are fully observed in the remainder of the transition phase and thereafter."

FOOTNOTE 3: SOME REFERENCES TO INTELLECTUAL PROPERTY ISSUES BY ICANN BOARD MEMBER JONATHON COHEN

According to an article "Net regulators close in on rogue domains" posted on April 10 at the Register by Graham Lea: http://www.theregister.co.uk/000410-000004.html

"Cybersquatters are facing defeat as a result of the Uniform Dispute Resolution Policy of the Internet Corporation for Assigned Names and Numbers. This was the conclusion to be drawn last week at the inaugural meeting of Cyberbe@t, a series of events for the intellectual property community arranged in London by Virtual Internet Net Searchers. Jonathan Cohen, a Canadian intellectual property lawyer and member of the ICANN board, drew attention to a number of problems that still need to be resolved. [Snip]

It has been a long road to get ICANN established and the UDRP working, Cohen told us. The objectors were generally from one of three groups: companies that had a vested interest in a particular outcome; objectors who claimed public policy motivations for their opposition; and dissidents who were known in some cases to have received financial support for attendance at meetings from NSI, as part of its effort to keep its monopoly. When questioned about the "constituencies" they claimed to represent, and the source of their funding to attend meetings, they were not forthcoming. It is now believed that there are new policies in place at NSI following the appointment of Roger Cochetti, formerly in charge of electronic commerce regulation at IBM, as senior VP and chief policy officer in February."

FOOTNOTE 4: FROM JOHN BERRYHILL TO WORKING GROUP B OF THE DNSO ON APRIL 14

(Referencing the Small Business Administration conference call on trademark issues in domain names earlier in the week of April 10, 2000.)

Mr. Menge,

Thank you for providing the opportunity for those concerned to present their views on the impact of proposed new gTLD provisions on small businesses, and for soliciting input into the upcoming SBA Office of Advocacy report. The following summarizes the two points that you had requested I send to you. Unfortunately, my legal training causes my "summaries" to sometimes be longer than the points themselves. As far as getting consensus on a compromise proposal, you can scroll to the Conclusion.

Point (1) - The Proposals Have No Basis In Technology Or Law

My comments essentially boil down to the fundamental maxim of Law, "Where there is a right, there is a remedy." The ICANN Intellectual Property Constituency's various exclusion or "sunrise" proposals are not in accordance with the remedial nature of the Law. These proposals are for prospective, pre-emptive restraints of the kind that we do not permit our own government to exert in the enforcement of criminal law relating to the use of words. Why should private individuals have greater power in the context of potential civil liability?

These proposals have perverted Law to "Where there is a right, there is a way to prevent people from violating it." That has never been the way Law functions in our society, and it has certainly never been the way the Internet functions. If it's not "technical administration", and if it is not "law", then I don't know what it is. Technical concerns say (a) domain name allocations are to follow RFC1591 - firs come, first served and (b) there is a need for a larger name space. The Law says that violations of private rights can be remedied after the fact. The IPC/WG-B proposals do not arise from valid technical or legal principles.

EricMengeIsaPedophile.com is libelous, and has legal consequences as a string of text. HaveSexWithMeForMoney.com is a criminal solicitation. TheHolocaustIsaJewishLie.com is likewise a criminal utterance, but in Germany, not the U.S. MuhammadTheProphetAtePork.com is blasphemous and likely a capital offense in several countries.

Yet, despite these and other categories of legally significant utterances, some even criminal in nature, nobody is proposing a prior restraint on them. Trademark infringement is only a subset of a much larger category of legally-proscribable uses of alphanumeric characters. Why, among all forms of legally significant text strings, are trademarks singled out for a heretofore unknown pre-emptive right? Because ICANN, a technical body, has an "Intellectual Property Constituency" with non-technical concerns. There is no "Libel Constituency", "Criminal Solicitation Constituency", or "Religious Constituency". Why not? Because these issues do not relate to technical administration, which is the mandated mission of ICANN. These are questions you can pass along to the GAO personnel studying the structural and delegated authority aspects of ICANN.

Despite the talk about the "importance of stability to the development of e-commerce", ICANN was not chartered to be about commerce or whatever else for which the internet might be used. They are supposed to be running narrow technical aspects of a computer network. "Do the bits get from one end of a wire to the other?" is not a legal question. Re-engineering the remedial principle of law as a proscriptive technical policy makes no sense.

Trademark infringement happens in telephone book listings. All kinds of shady folks get fradulent telephone book listings, or use "Yellow Page" ads which infringe trademarks or convey a false or unfair commercial impression. These situations are dealt with all of the time by trademark lawyers. They are not dealt with by providing a preemptive famous name list or a sunrise period for telephone books. In fact, the makers of the telephone books are not held liable for these kinds of things. In the context of 800 number assignments, the FCC has decided that dealing with trademark issues is a job for trademark lawyers, and not for technology policy makers at the FCC. Why should ICANN be any different?

The DNS is a telephone book. It maps names to numbers in precisely the same way. Why is it that we manage to publish telephone books without difficulty? Why would we argue about adding a new telephone exchange in an area code, become concerned that the possibility of a greater number of telephone listings would provide more opportunities for trademark infringement, and suggest that it would subject the telephone book publishers to legal liability? Because they are ridiculous assertions. But somehow they are taken seriously in the context of the DNS.

Even when someone has successfully asserted a trademark right involving a telephone listing, the books themselves are not published again until a year later. The DNS can be altered within a matter of hours to reflect a succesful, and remedial, assertion of trademark rights. That serves the interests of IP owners even more efficiently than an analogous system -phone books - with which we have lived comfortably for years.

To make the picture even clearer. I can infringe trademarks with my business card, letterhead stationery or outdoor signs. But when I walk into the print shop, there is no IP daemon sitting on the shoulder of the printer with the job of determining what words I may or may not have imprinted on my business materials. I bear the legal consequences of my choice, but I am as free as anyone else to have my own business materials without having to wait outside during a "sunrise period" in which the "first among equals" negotiated what is to be left over for me to have.

And so we develop a byzantine system of chartered and non-chartered TLDs, and a system of restrictions and lists and sunrise periods on top of that. The next day after I, a lowly individual, am allowed to register domain names with the great unwashed masses, I obtain generic.generic (in the new "generic" TLD). And the day after that I set up my server to resolve kodak.ibm.cocacola.generic.generic/kiddieporn.html . Then what did any of this nonsense buy for anyone other than delay and large expense account bills?

Bold prediction #1 - there will continue to be rampant intellectual property violations on the Internet. Bold prediction #2 - there will be no way to prevent it, but there will remain remedies at law.

Point (2) - Artificial Constriction of the Name Space by the IPC is Hurting Small Business

There already are mechanisms to enforce trademark rights in cyberspace - the UDRP and the ACPA among them. Both of these mechanisms are available to anyone who can afford a lawyer, which, with the UDRP includes many but not all small businesses. Genuine cybersquatting hurts small businesses in smaller gross monetary terms, but perhaps in larger proportionate terms for the affected businesses, than it does larger businesses.

However, when BigBusinessCo is faced with a squatter who has occupied BigBusinessCo.com, .net and .org, then BigBusinessCo can readily afford to get rid of the squatter. Joe's Fish Market is faced with a much larger problem, because they cannot so readily afford to do the same thing.

The presence of a large, and I mean very large, number of TLDs does two things to help Joe's Fish Market - it increases the cost of pre-emptive cybersquatting and it decreases the value of any one domain name occupied but not used.

If someone is sitting on the domain "cocacola.irrelevant", not producing any content at a corresponding website, and demanding thousands of dollars from Coca-Cola, then why would anyone, including Coca-Cola care? The commercial injury to Coca-Cola of a tiny vacant island in a sea of thousands of TLDs is approximately zero. In fact, it is actually zero. Now, yes, there is such a thing as trademark infringement, but if the only thing one sees at a web site is "This Domain for Sale!" or "We Registered At Lousynames.com!" then what is the basis for any consumer to be confused about anything? They were looking for a brown fuzzy beverage in a red can. "Hmm.... must not be at this domain name...."

Conclusion

You had floated the compromise proposal of a mixture of "chartered" versus "non-chartered" TLDs, and how many of each there should be. All I could think of during that portion of the discussion is to consider whether it would be a good idea to have a large quantity of even numbers or odd numbers. In fact, there is no good reason not to have an infinite supply of both.

The mechanisms for restricting registrations according to various pre-emptive systems are flawed technically as they do not accord with RFC1591, and they are flawed legally as they do not accord with the remedial character of Law as we in the West have come to know it over a learning curve of hundreds of years. The IPC does not have the technical background to run the Internet, and WG-B does not have the legal sophistication to re-write basic trademark law. This is not how to run a computer network.

Executive Summary

Akamai, pp. 1- 12

Rapid and reliable delivery of web-based content anywhere in the world has become one of the most critical issues in enabling the continued the scaling of the Internet. Web caching started out in 1996 as an attempt of many ISPs to store locally as much of the content of the web as possible. Each ISP would make its own decisions about what content to fetch and how often to do it. This system created many problems for web content providers because they had no knowledge about what was cached where, by whom and with what frequency. Furthermore, since caching distributed their content to many sites, they had no reliable way of reporting to their advertisers how many people had seen the material. It was a hit or miss system that no one was happy with and one that created a major opportunity for others to fill. A year ago Sandpiper and Akamai were the most talked about competitors. We note that since then Sandpiper has been acquired by Digital Island and has been focusing on the rapidly growing field of business-to-business e-commerce, leaving Akamai as the acknowledged leader in general content delivery.

In late 1999 Avi Freedman left his position as Vice President of Engineering for AboveNet to become Vice President of Network Architecture for Akamai. We publish a long interview where Avi explains in detail Akamai's extremely interestng business model. What Akamai does is enabled by a very significant new use for DNS that it has developed.

Akamai it has its own network of DNS servers that keep in contact with each other globally. Akamai's other servers take the web content of Akamai's customers and store it in hundreds and then thousands of copies at the edge of the network as Akamai's global network of servers continues to grow. Akamai solves the problem of the world wide wait by pushing content as close to the end user as possible.

Akamai's network of DNS servers then accomplishes a kind of global air traffic control task of communicating among themselves network traffic conditions in real time to determine which local server to send a user's request to or, in the event that regional traffic problems are interfering with local reachability, how to retrieve the data from a more distant server.

Within a site Akamai figures out what data is not constantly updated. That data is migrated to Akamai's edge servers on a regular basis. The minimum amount of data possible is pumped from the host web sites to the edges, while each edge web server is kept constantly informed of the best path to get to the fresh host data it needs. Akamai charges each web site owner for the aggregate amount of its data delivered to end users anywhere in the Internet. The table (at the end of the interview on page 12) shows how many networks receive what percent of Akamai's total aggregate of content traffic. Its intelligent overlay network of DNS servers that direct web content look up must keep very good statistics so that Akamai knows what to bill each of its customers who pay to have their web sites included in Akamai's distribution network.

Akamai has, in effect, created a virtual private overlay of the internet where, as much as possible, it keeps packets on a single network and minimizes their having to flow upstream to transit from one backbone to another (where most packet loss occurs) and them move to the downstream customers of the other backbone. This means that Akamai can go to an ISP and ask to place its servers in the its key POPs for no co-location charge and no charge for bandwidth used. Why? Because it can generally show every ISP how, with Akamai servers locally, its customers will pull far less web traffic across the ISP's backbone that they would if the ISP tried to do its own caching. Or, if the ISP just sent the packets back and forth to the content provider's central servers. In addition Akamai can demonstrate how, in return for nothing more than some co-lo space and bandwidth, the ISP will save bandwidth and give its customers better service.

Freedman also describes how Akamai must deal with the needs of its customer's central servers that are most often located at large web hosting centers at major backbone sites. In these cases he may act as an advocate for the Akamai customer in procuring if necessary some Akamai owned and operated short haul links to ensure that they can have enough burstable bandwidth to meet peak traffic periods. Given his experience at AboveNet which ran this type of operation, he is well equipped to deal with the web based, content provider, the web farm backbone operator and the large number of down stream networks where delivery oriented servers can be placed as close to customers as possible.

Akamai has taken advantage of a narrow window of opportunity to become, in contrast to the older generation vertically integrated backbones, one of a small but growing number of content distribution networks. Such a network hopes to solve problems like the peering problem for a BBN which in the summer of 1998 rebelled at granting Exodus free peering because Exodus dumped more traffic into BBN than it took out.

Commoditizing Bandwidth pp. 13-18

Focusing on her role in ongoing efforts to develop a commodity market in bandwidth, we interview Lin Franks of Andersen Consulting. Lin helps to bring a non Internet protocol perspective to the issue of bandwidth commoditization by explaining in some detail her role in the commoditization of oil, natural gas and electricity. In the mid 90s she went to work for Portland General Electric. When Enron acquired Portland in early 1998 she met Stan Hanks. The interview recounts how she worked with Hanks to learn the technology issues for trading bandwidth while techning the internet technical people what skills were necessary to successful commodities trading.

For the past year Franks has been working with Andersen consulting in developing a training program that will acquit executives at the large carriers with the issues behind bandwidth commoditization, make certain that they understand the staffing that must be done to get their companies ready to participate in bandwidth trading, to help them form an industry group that can agree on a benchmark and standard contract and to coach them through the process of carrying out the first trades. She notes that a nascent industry trading association had its first meeting in Washington, DC on March 23.

Although the price of bandwidth is declining, Franks is bullish on its future. She states that "the real disruptive event is the realization by those in the industry that, regardless of technological advances, regardless of their fever to lay fiber across the world, there is really no way that supply is going to consistently be able to keep up with the ungodly increase in demand at which we are looking. Whereas there may appear to be a supply glut right now, bandwidth demand will rise to fill it and will exceed the available supply. Then there will be another technological advance that will increase the available bandwidth supply. Then demand will rise again and so on and so forth."

IBGP in Network Design, pp. 18 - 19

From the inet-access list we present a brief discussion on the complexity of design decision issues involving the use of IBGP in network design.

ICANN Footnotes, pp. 20 - 22

We present a brief series of ICANN goings that readers won't be seeing in ICANN

Continued on page 24

Continued from page 23

press releases. The first is the announcement from last October of a new investigation of ICANN by the GAO - one that was announced with such stealth at the time that we missed it entirely and from which we expect very little.

The second is a much more interesting April 7th, 2000 report from the European Union noting with confusion that the Department of Commerce's public position on ICANN is that it is it (DoC) and not ICANN which is in control of the DNS and ROOT. The report complains that U.S. officials nevertheless are still promising to turn over all powers to ICANN by October 1, 2000. It concludes with the expectation that just before American elections in the fall all power will be turned over to an ICANN under the watchfull eye of the Government Advisory Committee. Overall it shows how raw European feelings have become on the subject of American Internet technology leadership.

The third is a brief from London on a presentation by ICANN Board Member Johnathon Cohen in which Cohen is alleged to have said: " It is now believed that there are new policies in place at NSI following the appointment of Roger Cochetti, formerly in charge of electronic commerce regulation at IBM, as senior VP and chief policy officer in February." Yet another IBM monetary contribution to ICANN is announced at http://www.icann.org/correspondence/ibm-letter-24sept99.htm.

Finally the fourth footnote is a brilliant April 14th critique by patent attorney John Berryhill. Berryhill shows the absurdity of the Michael Palage plan to grant trademark interests full sway over a list of famous marks and 20 variants for each one known as the Sunrise proposal. While Berryhill's comments need to be read in their entirety the following excerpt gives a good flavor. "My comments essentially boil down to the fundamental maxim of Law, 'Where there is a right, there is a remedy.' The ICANN Intellectual Property Constituency's various exclusion or 'sunrise' proposals are not in accordance with the remedial nature of the Law. These proposals are for prospective, pre-emptive restraints of the kind that we do not permit our own government to exert in the enforcement of criminal law relating to the use of words. Why should private individuals have greater power in the context of potential civil liability?"

We predict however that, effectively, the fix is in and despite widespread opposition the Board will do what it is told by its trademark masters.

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