



New Business Model for All IP Network Enron's IP Over Fiber Network Alternative Way To Deliver High Bandwidth Web Content Overlay of Public Internet Has Major QoS Implications

Editor's Note: Stan Hanks is Vice President, Research and Technology for Enron Communications. We interviewed Stan on February 10.

COOK Report: Let's talk a bit about your background. How did you get involved in the Internet?

Hanks: I got involved in the early days of TCP/IP in the ARPANet in the late '70s and early '80s. I was a consultant for MFS in the early days of their starting their DataNet oriented products. At that time, Scott Yeager in Houston was basically driving for the productization of some kind of data oriented services. In 1991, working with Scott, I conceived the idea of IP based Virtual Private Internet's (VPNs) for use in solving the problem faced by MFS in having many fully wired but otherwise isolated office buildings. I showed MFS how to connect these buildings using FDDI in the form of a metropolitan area network. This is where the GRE tunneling protocol and IP VPNs really came from.

From this point we got the idea of connecting the cities by means of a private line network. It was so difficult to get DS 3 IP routers in those days that we wound up using ATM for the interconnection effort. Now the next deal after this was hooking Rick Adams up with MFS and putting together a virtual private network backbone for UUNET across the MFS ATM network. This was how the Internet and MFS came together. The concept of backbone interconnection by way of MAE East also came out of my work with Rick Adams and Marty Schofstell of PSI net.

At the end of 92 I left MFS and, after some involvement in building corporate

telecommuting networks, I went to Bechtel to put together to Genuity. Unfortunately Bechtel did not understand the Internet and, as a result, I walked away from the job. They understood it only as something they could get an enormous stock price for. Rodney Joffe who started Genuity had been sold his original Internet connection through Los Nettos by Jon Postel. Jon became a kind of mentor and educator for Rodney and wound up on the Genuity board. Rodney and another South African friend on the Bechtel board saw the possibilities of the Web and took Rodney's small Internet oriented business and spun it up into Genuity. They took Jon along through the path as they did it. Jon's influence was important in setting the right tone and in getting the right people involved. For example he brought Paul Vixie in as a consultant in the early days. As to the propriety of Jon being on the board, I think it never occurred to the people involved that the role of senior adviser to the board rather than a board member would have been a cleaner way to accomplish the end.

Around the time that I finished my work at Genuity, I get a call from a head hunter. The pitch was that a multi national company was building a network and was going to do it right this time. My first response was that I wasn't interested. They kept coming back to me and when, after a week, they finally told me the company was Enron, it took me about two hours to schedule a trip to the Portland office. I had worked with Enron in the past and, having lived in Houston, I knew an awful lot about them. At that time, for a third year in a row, they had just been tagged by Fortune magazine as the most innovative company in America.

Volume VIII, No. 1, April 1999
ISSN 1071 - 6327

COOK Report: I know a bit about Schlumberger. They are an energy company with a big data network. Is there any similarity between the two companies?

Enron Corp's Knowledge Oriented Market Position

Hanks: No. Absolutely none. Enron *does* have a large robust enterprise network, it's just that we (Enron Communications) don't have anything to do with it. I'm not actually sure how big the Enron enterprise network is, but it does cover operations in 167 countries. Enron communications is what you would call a pure carrier play. We are not doing any communications for our parent company.

COOK Report: So Enron Corp. buys its data communications services on the open market? Have they bought any services from Enron Communications yet?

Hanks: As of yet they have not. I think that over time we will probably evolve

On the Inside:

Enron All IP Net	pp 1 - 10
IP Telephony Prot Design	pp. 11 - 15
Nepali ISP	pp. 16 - 19
ITU and Standards	pp. 20 - 21
ICANN Wars	pp.19, 21, 22, 24

some services that they will use. However selling services to our parent company is by no means the reason for our existence. If you stop and look at the history of Enron Corp., you will see that Enron has grown its business by identifying and exploiting opportunities that arise from market dislocations. It makes investments in the outcome side of the infrastructure rather than in the incumbent side. This gives it a tremendous advantage relative to its competition. They have done this in oil, and natural gas, and in electricity.

COOK Report: Would arbitrage encompass a portion of what they do?

Hanks: Arbitrage is a significant portion of what they do. If you look at the 10k for Enron Corp. you will find that a significant percentage of the company's net income comes from this small organization they have called Enron Capital and Trade which is basically a huge commodity trading house.

COOK Report: so would you say that Enron Corp. functions as a company that is involved not so much in energy raw materials as in information and knowledge management within a broad scope of energy industries? And furthermore that they use their knowledge to try to ascertain where new business opportunities are emerging and to enter the markets created by those opportunities before anyone else?

Hanks: I'd say that's pretty fair assessment. They tend also to work a lot smarter and in a lot more capital efficient manner than many of their competitors. For example, in natural gas, without owning a big gas distribution network, they managed to rise to the top as the biggest buyer and seller of natural gas. They did this by figuring out other economic ways to make the market work.

COOK Report: So how then did you come to fit into their picture?

Hanks: A few years ago Portland General Electric, one of the power companies in the Northwest, had been approached on numerous occasions with the idea of building fiber-optic networks along their rights of way. They always declined the opportunity. The genesis of Electric Lightwave and of GST Telecom both came from opportunities that Portland General Electric decided to pass on.

COOK Report: Are we talking about including fiber cables as a part of the ground wire that is strung from the top of one high tension tower to the next?

Hanks: yes. This is generally the best way to along power line rights of way. The fiber is usually either wrapped around the ground wire or is part of the ground wire bundle itself.

COOK Report: Is it pretty cheap to install?

Hanks: Well, relative to holes in the ground, it is fairly inexpensive. But remember that you need things like helicopter crews to do it. The good news is that you don't have to dig any holes in the ground and the bad news is that in parts of the rural West you have to worry about young men shooting your cable off the towers just to prove that they can. Nevertheless stringing fiber between high tension towers is a fairly localized phenomenon. We're not doing a whole lot of it.

Power Company Network Infrastructure

In 1996 Portland General Electric, having missed out on several earlier opportunities, engaged some consultants to examine current market alternatives. They concluded that they should take their deregulated company structure and use it to start a new company which, at that point, was called First Point Communications. Its initial mission was to use Portland General's right-of-ways and relationships to build large scale fiber-optic distribution networks within the Northwest.

As part of this effort they wound up in a consortium kicking off construction with Touch America which is Montana Power and with Vivix, which was and still is, Williams, to build a fiber run from Portland to Los Angeles going by way of Boise, Salt Lake City and Las Vegas rather than down Interstate 5 along the Pacific Coast. This reverse crescent route on the eastern side of the Sierras was tremendously valuable for all the other fiber routes basically do go down Interstate 5. This route gives people a diverse path should anything bad happen along I - 5.

COOK Report Backup potential in other words?

Hanks: Yes. And about the time all this was kicking off in 1997, Enron bought Portland General. Enron had watched to the Williams people up in Oklahoma make just a ton of money from selling Willtel in 1994. Thinking about this after the purchase, they looked at the Portland

General situation and said, you know we ought to be in the communications business also. By buying Portland General Electric in the first place Enron was improving its ability to compete in the electric power markets of the Northwest. They had been trying to encourage the deregulation electricity on a nationwide basis. By controlling a significant portion of the distribution capability in the Pacific Northwest, they wherein a position to influence at the very least what happens in the principal part of the community

At this point, a 1997, they realized that they needed to do something with IP and the Internet but they really didn't know what. Because of the people whom they had associated with their project at this time, they were mostly thinking of CLEC oriented things. However Enron thinks globally not locally. As a result, over the course of a few months, we redefined the mission of the company and instituted some management changes.

The reason for the global focus was the market dislocation effect. If I can, via technology, make location be completely irrelevant, this opens up whole new ways of working and hence, entire markets. Of course, an added benefit is when you go into developing countries and are pitching big infrastructure projects like highways and hydroelectric construction and you can also say: oh, by the way, we have this communications group that can bring you global telecommunications infrastructure as well. After all, if we're opening the opening the ground up for other construction purposes, all we have to do is throw in some fiber-optic cable to give you a state-of-the-art telecommunications network as well. Of course most of the money that is to be made is in North America and in Europe these days. However we do not think that this will always be the case — especially if some of regulatory problems in the Third World can be solved.

Choosing the Right Technology

So we looked at things and saw that it would take about three years to build a network, therefore we must now decide what kind of data network we want to build. Will it be the network that you wished you had three years ago or the kind of network that you would really like to have three years from now? The decision to build Pure IP basically came from the fact that IP is the convergence layer. IP is where new investment, mind share and

general technology development effort is directed.

COOK Report: So when you picked this direction how did you position yourself vis-a-vis the earlier arrivals of Qwest and Level 3?

Hanks: You touch upon a question into which we put a fair amount of thought. Now Qwest was saying that they were going to be the low-cost IP provider. However, added that time, you would look their Web page and what was the first thing you saw? Telephone calling cards. Level 3 was saying that they were going to build this all-new IP end-to-end network. But, if you look at the equipment they were buying, you will find that they were basically re-creating the old MCI network.

COOK Report: And if you look their IPDC protocol development and the associated developments that followed, you will see that its purpose of is to enable the carriage and overlay of the public switched telephone network on an IP network.

Hanks: Absolutely correct. The CORE of what they have going on the transmission side is all ATM-based with frame relay for aggregation and with IP kind of in the middle of it. It kind of bothered me that they were saying IP this and IP that because when you really pushed into what they offered, it "gave" a lot. The IP in this case was really there so that they could sell Internet services. All this is familiar turf for me. It represents also the battles I was fighting with MFS in the early '90s as well because the regular phone companies increasingly pushed back on all our service offerings saying they would rather do them on some TDM infrastructure.

COOK Report: Would you disagree then with the evaluation of Level 3 as Jim Crow's effort to replay Metropolitan Fiber Systems infrastructure in terms of the latest technology?

Hanks: Jim Crow has dial tone envy. If you stop and listen to the interviews with Jim and realize what he is saying, then you will come to understand that he wants to be AT&T before the breakup, but he wants to do this using technology that is available today so that, instead of investing a hundred billion dollars, he invests between 4 and 5 billion. He wants to carry all the bits he can. He wants voice. He wants video. He wants Internet. If you are a residential end-user of consumer services, he wants to provide them all to you.

COOK Report: So are you saying of that all these traditional service markets are rather crowded and asking why not leave them to the people who are already there and go off in a direction that would give you some possibilities and strengths that these other people could not easily emulate?

Hanks: Yes, this is essentially correct. For Enron it was not terribly difficult judgment call. They could reasonably conclude that, before long, the only thing that someone would wish to transmit will be IP. Therefore if we build a network could only transmits IP, we're making a safe bet. If we structure the building process correctly, we will have the latitude of one of the most important elements of the Qwest model which is financing our expansion by selling the access to our dark fiber. Therefore we construct a network that pays for itself as it goes by selling dark fiber and we don't have to worry about whether not we can pull in revenue from every conceivable telecommunications sector.

COOK Report: So are you saying that by selling access to the reverse crescent of Portland to LA infrastructure than you have laid, you gain the ability to go to a company like Frontier and offer them access to your infrastructure in return for your access to theirs?

Hanks: What we'd did with Frontier was precisely that — a trade of capacity on our network for capacity on their network. I suppose that you could also look at it as two purchases that have approximately the same value. Our Web pages list several fiber swaps that we have done with companies like Electric Lightwave, GST and Fonorola. The GST piece goes up Interstate 5 from LA to Portland giving us a ring on the West coast. We have also undertaken construction on a piece of fiber from Denver to Houston and are about 30 percent done with that route. That route again is something we will sell dark fiber on and swap capacity for.

COOK Report: When you get your spreadsheets going, to what extent do you find that the this build, buy, and barter modus operandi is a self-financing model?

Hanks: It comes pretty close. Suffice it to say that Enron Communications has been profitable for the last two years.

COOK Report: Are you contemplating an IPO?

Hanks: Enron Corporation has spun off

several units before. And, given the ridiculously high values that networking and communications companies are enjoying right now, I think it is a foregone conclusion that Enron will want to capitalize on these trends. Whether they will do it by spinning us off in an IPO while they hold the majority of the ownership of the company still and enjoying the value inflation on their balance sheets is a result of that., I honestly don't know. Given current economic conditions I am sure there are as a whole team of people in Houston sweating over the outcome. However there is no offering of which I am aware that is definitely scheduled to take place this year.

COOK Report: Well, so much then for your infrastructure build out. How did you derive your business model?


The Enron Communications Business Model

Hanks: Enron has a market experience that really influenced us. Retail sales can be really messy because you are dealing with large numbers of consumers who are buying medium to small sized amounts of service. You spend a large amount of money and servicing the sale, customer billing and so on. If you go retail therefore, you either have to get to tremendous market penetration in order to become profitable, or you have to focus your attention in niche markets. This retail structure simply wouldn't work well for us. Consequently we adopted a position that has us deeply embedded in wholesale markets.

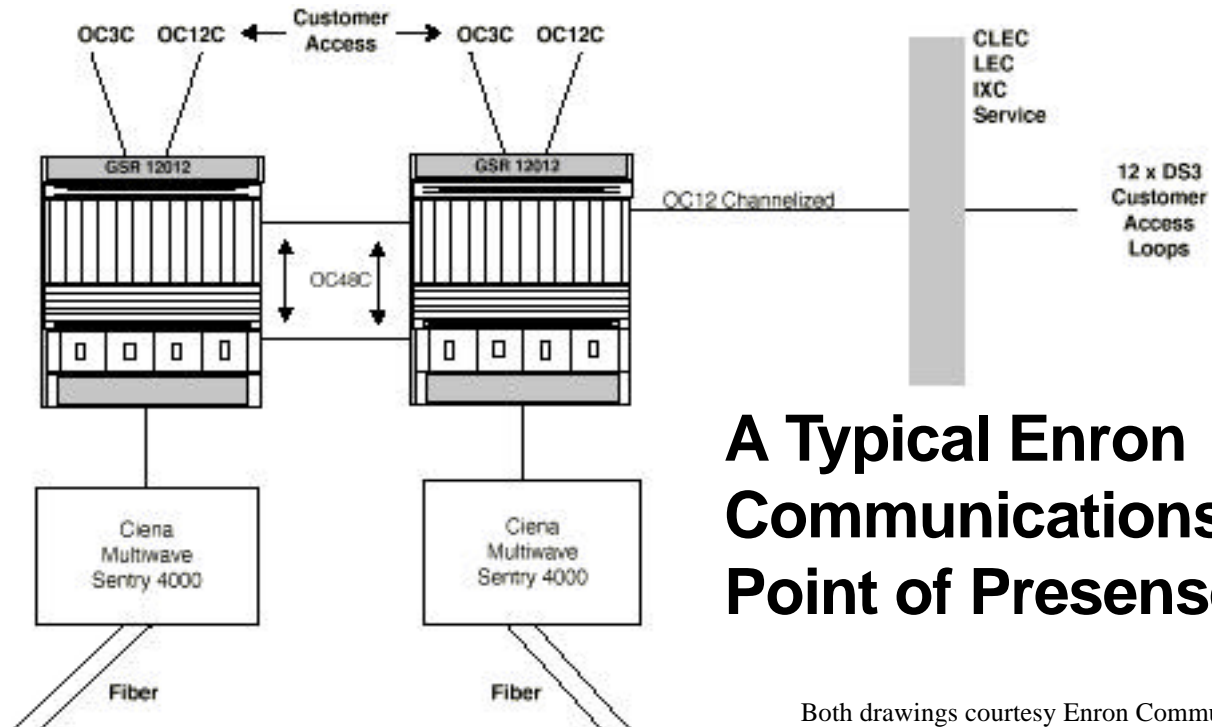
The term carrier's carrier is used to describe us quite often. However, we don't much like this term because, to us, it connotes a telephony mindset.

We have an all IP transmission facility and an intelligent network structure that we also need to talk about. The simplicity of our own POP architecture is illustrated by the two diagrams on the next page.

We enable people who are engaged in local information distribution and content aggregation to do things that they would be unable to do otherwise. Rather than competing against our channel partners, we have a revenue sharing model so that, rather than seeing us as competition, local providers find us to be a source of income. In order to elaborate on this I have to explain our intelligent network structure.



A Typical "Industry Standard" Point of Presense



A Typical Enron Communications Point of Presense

Both drawings courtesy Enron Communications

When you stop and look at why people use networks, you will find that no one really cares about pushing bits across the wire but rather that they care about what happens as a result of the bits that are pushed. What matters to them is that they made a phone call, downloaded a Web page, or saw a movie. If you look at the intermediate piece between them and the end result you will always find that there is a computer in the middle. It may be converting an analog signal into digital. Or it may be converting an image on a screen into a bit stream that across a network. What you are dealing with is really a distributed computing network and not a data transmission network. In our model the computation that takes place within the network is a critical part of the information delivery.

Now when you look at the Web caching model, you will discover that the reason that it exists and that why networks are built the way they are, is that people were unable to buy a service having a computer-based service point of demarcation between the network and the end user.

If you look at a traditional telephony carrier, your service demarc is a terminal on a multiplexer. At each end of the connection you must make local connections across another network that looks like the first in order to get to the point where you have a digital pathway to your premise where you can place your routers and your interconnect LAN to your local computers on the local desk tops. You are having to deal with a lot of local complexity because of where the service boundary between your network and that of the carrier lies.

Now if you look at an ISP, you will find that it moves the point of interconnection (the service demarc) up a little bit to the point where it is now an interface on a router. They handle some of the local loop problems and some of routing problems for you, but you still have full responsibility for local distribution and for your computing core.

Contrast this with our business model where our service demarc is a TCP port on our server. If you look at the initial service deployment that we have announced with Real Networks, you will find that we are providing very high-speed access to servers on a local geography basis. We are providing streaming media services at high bit rates. Now today when you log onto the Internet and point your Web browser to CNN, what you get is a TCP bit stream all the way from your

machine back to the main CNN Web server. In our network model when you asked for the same thing, what you would get is a TCP stream going only to our local server and not all the way across the Internet. We handle, through replication and distributed computing technology, the distribution of that content to the edge point of the network closest to you. So you could say it is form of caching for real-time video and audio. It's not quite the same thing, but the principles behind both are similar. It's a bit like caching on a sub micro-second basis. Now we have other services that we have not announced yet based around the notion of what you can do once you have decided that a significant portion of your distribution infrastructure is actually made up of computing resources.

COOK Report: Are you saying that you can use your network to tie together a network for them where in major areas they can have local pops so that, if a content feed originates from Seattle, it can go through the Real Networks pop there and be extremely quickly delivered to every other pop in your network traveling free of the rest of the Internet.

Hanks: Yes. If you start looking at what we are actually doing here, Real Networks is a content aggregator in our business model. Therefore we have a high-speed connection into their distribution center in Seattle. The content that they originate and we distribute gets sent instantaneously to all of our servers at all of our network pops. From those pop's we have high-speed connections to ISPs, CLEC's and cable modem people.

COOK Report: So you become a distribution channel for Real Networks in such a way that you can get their content to local infrastructure providers and hence to the audiences they want to reach in those areas more efficiently and cost effectively through your network than through the public Internet?

Hanks: Correct. And the other piece that goes with this is that we offer a managed infrastructure where we handle all the plumbing and all the computing aspects of delivering the content. For example consider the local distribution partner. We drop a router and a server under our control and management into their pop. Along with this we pull, at our expense, a high-speed leased line into their pop. We then provide from our infrastructure to theirs whatever high-speed interfaces are necessary — typically multiple 100baseTX connections.

Choosing Local Partners

COOK Report: Let's follow the import of this a little further. What if your local distribution partner were a regional ISP like a Net Access in Philadelphia? The owner, Avi Friedman, has about 10 local ISP's, 50 leased line customers and about 6000 dial-up customers hanging off his infrastructure. Does an ISP like this fit into your market? How would you approach such an ISP?

Hanks: Certainly he would be a potential fit. We would bring in the leased line, the router, and the server at our expense. There would be some agreements that he would have to sign in terms of meeting service levels to his customers. The criteria for these agreements will be determined in accordance with the criteria called for in our agreement with Real Networks on the originating end. Assume that he has met all the criteria in that everything has been installed. We will actually pay him money for each of his subscribers who gets a Real Networks feeds through our infrastructure.

Let me make very clear that this is not our entire business model. For this particular product, this is the way the model works. Real Networks is paid by the people who push content into them to carry that content. Real pays us to carry their content based on how many people view it. We, in turn, pay the local ISPs and CLEC's to distribute the content to the end viewer.

Now note that I have put a T3 link into his pop at my expense and I pay him for some of the services that go over that link. However I may also be able to sell him other services across that same link. Note all so that the cost of those additional services might well be paid for by the income that he gets from delivering the Real Networks traffic.

COOK Report: Well what might those services be?

Hanks: We are not talking about those other services just yet. But, if you think about this, you can see applications out there like videoconferencing and collaborative workgroup computing. We think we can deliver high enough bandwidth cheaply enough to make the premise that the computer is the network come true.

COOK Report: Is there a critical mass of distribution channels that you have to achieve before you can think of opening

up a real-time video conferencing service? Will you be saying to the CLEC's and the ISPs in the cities you serve: we have a videoconferencing service available that plugs-in to when you wish for such and such a price.

Hanks: Yes, that's one aspect of it. As far as critical mass is concerned, you can basically say that it begins to appear when you have enough bandwidth between enough locations that matter. If you start to open your pops in major U.S. business centers, as you go down the list, you will accumulate critical mass rather rapidly.

It also depends on what your objective is. If you have a push oriented content play that works off the revenue model that I explained in the case of Real Networks, then you want have as many pairs of eyes staring at your screens as you possibly can get.

COOK Report: But this is not really the kind of push model that was all the rage a couple of years ago because this involves actual customer choice?"

Hanks: "True. It is really something close to a cable TV model because it says select your information channel of choice. But note that you will also be able to vary the capacity of that channel with respect to the quality of the image that you desire. You can take a feed optimized for a hundred Kbps or for a Mbps. But keep in mind that we are primarily planning to use this to deliver business information content such as corporate training classes, sales meetings and that sort of thing.

Handling Traffic Flows Creatively

To summarize: we bring a very high-speed link into the local ISP's pop and some amount of that high-speed capacity goes for delivering certain services. When you stop and think about it because you are staging the content from the server at their pop, you can have 10,000 users hitting the server at half a megabit each for a total of five gigabits of outflow but I still have to have a only half a megabit inbound because it is all getting staged off the local server.

COOK Report: But figures like these would require huge local infrastructure?

Hanks: That's true. We think our prime candidates for partners are those who are delivering a Mbps to their customers anyway. Namely our primary local

partners look like those who have a high concentration of T1s into corporate customers, or are heavy DSL players, or are into cable modems and a big way. So if you ask the question: as bad as dial-up sucks going to my ISP, what's a DSL going to be like it going to the same guy? We can be the answer to that same issue. To achieve this we provide a crafted infrastructure designed to deliver outflows huge amounts of information.

COOK Report: Are you arguing that, they get the full range of content data and ISP more effectively and more cheaply than would be possible if the ISP were multihomed to Sprint, BBN and UUNET?

Hanks: Yes that is correct. The bottom line is that we put all this information one hop away from his network's core as opposed to having to travel so many hops through my backbone, a couple of MAEs in between.

Peering Is Not a Part Of the Business Model

COOK Report: But what does this do to the whole concept of peering?

Hanks: Here is where life gets real interesting because we have peering meaning two different things. Let's consider the technical meaning of peering which is a technology implemented by means of BGP4. It lets me exchange selected routing information with another network. End of story.

COOK Report from a purely technical point of view we can turn up BGP4 peering where of my routes go freely into your network in your networks routes go freely into mine. Right?

Hanks: Absolutely. And, even more to the point, that kind of stuff winds up being absolutely essential in all the kinds of connections that we have been talking about. Now the other definition of peering is known as free peering for the purpose of providing transit. We don't do that. You cannot buy transit from us. We don't sell it. It's not something we do. In this sense we remove ourselves from the peering wars. If, on the other hand, you have an autonomous system that happens to be the Real Networks broadcast network and that you want to be able to deliver content into your network, about the only way that you could accomplish this, frankly, is to use some kind of BGP based peering technology.

Remember that in addition to the we-pay-you-to-watch-content model, we have this other model that says you pay to watch specific content that you're interested in. We have to be able to discriminate on such traffic on an origin basis. Consequently, in our network, traffic flows according to how it has been contracted for.

COOK Report: You achieve this by what means?

Hanks: On the technical network engineering side of things the means includes route filtering and having the right kind of announcements as to where stuff comes from and where it goes to. But where it really gets interesting is on the business side of things. In the Internet the assumption is, that if you inject a packet into the network, it will get to its eventual destination. However, packets in our network only get sent to contracted destinations. So, if you're not contracted to receive certain services, you simply will not get them. We do not replace an ISP's link into the Internet. We're not an upstream feed.

What we might do is divert a lot of that ISP's upstream traffic. Let's take this five gigabits scenario. You have 10,000 users hitting a local ISP at half a Mbps for Web downloads. Under the existing scenario, what would have been the case is that he would have to back haul all five gigabits across his upstream link across its backbone and through the MAEs or private interconnects to Real Networks. This is simply makes no sense, right?

COOK Report: Right. One of the reasons that ISP's like do as much peering as possible at as many different locations as possible this to keep their traffic as local as possible and to keep their major backbone uplinks free to handle non-local traffic. Therefore, taking a feed from you, would take that traffic away from the upstream traffic on which they would otherwise have to pay for transit?

Hanks: Right. Texas has a great example. Consider what happens when you want to get from Texas A and M University to Rice University. The two campuses in the Houston metropolitan area are about 75 miles from each other. However, to get traffic from one campus to the other, you have to send it through San Jose CA. What you must do then in order to really make this work, is to find the places in the Internet that originate the high-density content. This is the stuff that people typically pay a premium to get to or that needs high bandwidth. With our network,

you bypass the bulk of the Internet for rapid delivery to points of local distribution.

COOK Report: Do you remember the Exodus/BBN peering dispute from last fall? Does your business model have an answer here to problems Exodus, or to BBN or to both?

Hanks: I think we have an answer to both. To BBN I look like something that delivers high-speed content to their local distribution pops. And both the part of BBN that now runs Genuity, that runs data centers in other words, and Exodus both look like content aggregators in our model. We certainly have a potential fix for many of the things that are broken about the current Internet traffic scenario. The fix of course is to move the more demanding content across our transmission infrastructure. But also remember we do not provide general transit. You cannot point default at us. We are not a General backbone alternative. But when we show up, we may bring you a grand total of 17 high traffic addresses that you can get to at one hop as opposed to having to go through the entire mess of the Internet.

Depending and what your traffic load is, and on what your access problems are, and depending on how the economics of this model scale, you may find less and less of your egress bits going through your upstream provider. This changing traffic pattern will increase the bandwidth available for your upstream growth. We think this is the Internet economic equivalent of the discovery of fire.

What makes this possible is our intelligent network which we have only really been talking about since late January. We have been building this very quietly for some two years now. Software accounting mechanisms are part of it. The most important part, however, is something we call the control plane for the network. We are building this control plane using a piece of software called InterAgent. It was developed by a company called Modulus Technologies which we bought in November. This software was really developed for Robotics real-time control applications at NASA. It is a very powerful tool for use in building distributed control infrastructures.

COOK Report: So how would this software work with a Real Networks type of application?

Hanks: suppose you click a URL that says

you want to watch CNN interactive. Your request would go through the normal Internet control mechanisms and then hit the server at Real Networks. Then, based on the source address of your HTTP request, the server says "oh yeah, we know about these guys through the Enron network.." It then fires off a request mechanism into our control network which enables a flow to be set up going back to the requestor through our network and giving us credit with Real for the transaction. If Real can get to the requestor by means of the Enron Network, it sends the packets accordingly. But if a requestor is not reachable by means of our network, the real server sends the packets through the public Internet. For accounting purposes it aggregates the totals over time.

The Business Model from Enron's Point of View:

COOK Report: How easily copied is your business model?

Hanks: That's difficult to measure. In some respects, one could argue that it's not that difficult at all. On the other hand, I think that our financial model fundamentally breaks the core telecom pricing mindset, which basically is "pay for 100% of the facility that you use, despite the fact that you don't use it 100% of the time." And despite the "coopetition" buzz word that's tossed around — implying a blend of cooperation and competition" — you see that most of the telecom players out there still seek to own facilities end-to-end, instead of promoting cooperative deals. And god forbid that they share customer info, and certainly not revenue! Measured that way, I'd be surprised if to see anyone else doing this on our scale anytime soon.

COOK Report: How unique is the InterAgent software?

Hanks: It's fairly unique. I started life doing distributed systems work, and in The gap between MFS and Genuity, had a joint venture with IBM to productize some work from their research labs. I got pretty familiar with the various software out there in the distributed object middleware space. The InterAgent software is fairly unique, and there really isn't much out there that truly competes with it.

There's a lot more about the InterAgent software on their web page, at <http://www.modulus.com>.

www.modulus.com. We're still treating it as a product, though, and are actively licensing it to a variety of partners. This places us in a very different camp with respect to our core vendors. Many times, we are not only consumers of their technology, but now they're consumers of our technology as well.

COOK Report: What about some cost comparison studies of content providers working with you rather than on their own via traditional approaches.

Hanks: Well, let's go back to the two underlying aspects of life in the Internet. First, there's the telecom carrier model that prices transmission around the fact that since everything is TDM (time-division multiplexed). If I as a carrier dedicate transmission capacity to a customer between point A and point Z, that all of that capacity is "burned" for the benefit of this one customer, and that I can't recover any of the underused or idle capacity. So carrier prices tend to be fixed, and tend to apply regardless of how well utilized a circuit might be. This tends to discourage customers from buying enough network capacity to meet their maximum needs, because that might only happen a few times a day — but they *pay* for the capacity 100% of the time, regardless of utilization.

Second, there's the notion that all packets are to be routed equally, and that it's perfectly acceptable to heave large percentages of traffic across any two points on the network in peak utilization periods as long as you're really, really trying to route all of them. In part, this arises from the fact that in the Internet, you have no idea from moment to moment exactly what network origin and destination addresses will be trying to connect through your routing fabric. The notion of being able to say "hey, these bits didn't pay to come across here, so out they go!" is a bit vague when you talk to most Internet service providers.

So back to the cost comparisons: first, we're providing service on a utility basis, where you pay for what getting the bits there is worth rather than the cost of the underlying facilities. Second, we know exactly at all times what traffic *should* be in the network. It's a pure "no tickee, no washee" situation — if we don't have a contract for a certain type of service with you, we just don't carry the bits.

I believe that this puts our services in a highly advantageous position. Let's take

the Real Networks deal for a minute. At our cost, we place the high-speed DS-3 line into a distribution partner's POP. We also place the router there, and a local server. And we wind up sharing revenue with them for content distributed to their customers. One of our salesmen summed it up as "nothing down, we supply everything, and give you free money. what's not to like?"

COOK Report: Is there something about your network that supports real time audio or video in a way that it cannot be serviced otherwise?

Hanks: The underlying architecture of our network that depends on distributed computing and storage within the network as part of the fabric means that it's possible to do many, many things that can't be supported otherwise. For example, let's take a streamed video service. In our network, end-user clients connect to a server at the distribution POP, which in turns connects across our backbone. That means that we can accumulate the accounting info on the server, as opposed to trying to fish it out of the mass of data in the backbone accounting. Ask pretty much anyone in the industry how hard it is to get useful billing info from the netflow accounting records — they'll tell you that you can pretty much not see any finer granularity than per-address-space as opposed to per TCP connection, and that you really can't do this at speeds higher than about 100 Mbps reliably.

COOK Report: When will you have an Internet telephony customer? Do you need anything unique to support such a customer(s)?

Hanks: We're not quite ready to announce anything there just yet. As far as what we need to support such customers, that sort of depends on the customer requirements. To date, most of the large-scale VoIP gear makes assumptions about latency and packet loss that are pretty optimistic for use in the Internet. In our network, however, where we know all the flows in advance based on contracts, well, I'd expect things to be significantly simpler. We'll be talking more about this by mid-year.

COOK Report Do you have a strategy for using your bandwidth in place of quality of service?

Hanks: I've been on record for a long, long time as saying that until all the knobs and twiddles are in place to make IP QoS really work in a seamless and fully

interoperable way, that the one thing you can always do is kill the problems with bandwidth. After all, QoS isn't really relevant until you hit congestion in the network, right? So if you can jack up the bandwidth until you never ever see packet loss or congestion, then no matter how the packets are labeled, they all *effectively* have the same class of service, and everything "just works".

COOK Report: Dave van Allen, an owner of FASTNET, a major eastern regional ISP had the following to say about the first draft of this article: "Stan appears to have taken the economic strategy of QoS to a logical level; perform QoS in a technically heterogeneous if-then environment. The difficulty in implementing QoS is that many people want to carve out pieces of the core for QoS instead of overlying QoS on top of the core. Enron's idea is elegant, and if implemented in a fashion where the ISP POPs can be enabled in number and quality, is a solid step toward high quality of service content delivery. The model should be very interesting to content aggregators and very beneficial to multi-megabit end users, especially those who have a great or important need for the delivered content. In a word or three - I like it."

A Fresh Approach to Quality of Service

Let's explore the QoS issues further. Your approach is refreshingly different from every thing else we have seen.

Hanks: It all goes back to what it takes to deliver Quality of Service. This includes the IP class of service parameters where you have a very small number of bits to play with. You have to deal with the dynamic tension between the requirements of operating a core network and the requirements of operating the intelligent network at the edge. You find that, at the edge, if you want to, you may offer very many flavors or gradations of Service. But, at the same time you will find that as you go across your backbone, the limits of technology require you to cut down those flavors and gradations to very small number. In the IETF this shows up in the difference between the Differentiated Services and the Integrated Services working groups where they basically recognize that they can only deliver a small number of classes of service.

COOK Report: So we have to take the variety found at the edges and meld it all together in such a way that the backbone

routers at the core can handle it in the incredibly short amount of time they have to look at each packet as it flashes through.

Hanks: Right. At OC-48 a 40 byte packet transits the router in 200 nano seconds and in about 50 nano seconds at OC-192. There's really not much time to do anything with the packet besides forwarding it.

COOK Report: If you look at the Cisco 12,000, the biggest work horse of the Cisco line, what do they do in hardware with this router that would give them much of any kind of ability to do any tweaking of Quality of Service at the backbone? An how might you compare what Cisco can do in comparison with the data hardware architecture that VJ Kumar talks about with regard to the Lucent PacketStar router?

Hanks: The Cisco 12,000 has a set number of hardware queues per interface. It also has the concept of virtual queues so that you map every queue on one input edge to every queue on an output edge. Therefore you go from eight physical queues to 64 virtual queues per input output pair. Cisco manages most of that in hardware. It is certainly clear that they're pushing more functionality into the hardware end and doing less and less in software.

COOK Report: And just by looking at the increasing speeds and the laws of physics, it becomes obvious that one is virtually compelled to move things into hardware.

Hanks: Right. that means that one also must make trade-offs in the design phase in terms of picking the right things to be used in the classification of packets into flows and into classes of service. Once this is done, then you must also have a good way of tagging the packets so that you can send them in the direction that you want them to go. MPLS adds a great deal to this in terms of reducing the complexity involved. If you can edge mark it as it comes into your network and drop in into an MPLS flow, then you can basically say anything that is in this MPLS flow I know, a priori, has got these kinds of characteristics. The hard part comes when you have multiple different flows that you are trying to map together because you run into a kind of aggregation problem.

COOK Report: So when do you start mapping the flows together?

Hanks: I think you want to start as soon as you recognize that you have

commonality. For example the concept of a trunk flow is useful. It says that you go from a given address to a given address at approximately the same rate of service. Thinking along these lines means taking all of the smaller flows and aggregating them together. But that really only works if you do it as soon as possible. What you then wind up having are conditions in which these trunk flows merge and split as you get deeper into the network. Managing the whole process becomes pretty complex.

COOK Report: Some of these efforts sound almost connection oriented. But how can you mix in connection orientation with statistical multiplexing?

Hanks: Even with connection orientation you would have some of these problems because, when you engineer your circuit, you have to make sure that there are adequate resources available to cover every link of the path.

COOK Report: Well to put it another way, I have the impression that to get the Quality we want, we have to go back and pick up other aspects of a connection oriented architecture. Is there something to that idea?

Hanks: Yes. If you stop and examine what people are doing to traffic engineer their large backbones today, you will find that very many of them are doing their performance engineering and traffic management by means of virtual circuits through ATM fabrics. They simply engineer the virtual circuit to be an appropriate size with an appropriate class of service between two points.

Now when you stop and think about it, IP as we've traditionally deployed it, always looks like routers connected by point-to-point links or routers connected by buses. And what we know today about how to engineer the network comes basically from those sorts of concepts. The problems that we must solving in building big and fast Quality of Service networks that have nothing but IP in them are not trivial.

COOK Report: So how are you moving to do this in your network?

Hanks: First our intelligent network architecture basically minimizes the load that we have to push across our backbone. We are doing all of our fan out from computation and storage located at the edges of our network. This is a key element. The second piece is that Quality

of Service only becomes an issue in the presence of congestion in the network. Of course for some period of time, anyone who's building a really big network can, as one of my engineers likes to say, kill the problem with bandwidth. That means basically when you detect congestion, rather than figuring out how to twiddle your quality of service parameters, you just increase the bandwidth on that segment of the network.

COOK Report: But on the other hand, if you have something like the Packet Star router that can do a lot of QoS in hardware and you stuck it in the core of your network, it would be helpful to you, would it not?

Hanks: Yes it can but you still must first figure out how you are going to aggregate varying qualities of Service the edges of the network and put them into common trunks going across the network. You can, of course, solve this with a commercial decision by saying I have eight hardware queues, that means I have a grand total of eight classes of Service and here's what they are. You either fit in one of these or you don't. The hard part comes when you tried to derive general solutions that do everything you could conceivably want for every one. Then it just all starts breaking.

Limits to a Bandwidth Only QoS Solution

COOK Report: While you have relatively few customers, you can certainly apply a bandwidth solution to them all. But when you to grow from dozens to hundreds, to ultimately perhaps thousands of customers, your ability to throw bandwidth at problems will become much more elusive and you really begin to need network engineering solutions.

Hanks: Yes, and this speaks to two situations. One is that this is a service optimized network. It is designed to do a couple of things very very well. As it becomes possible to put additional capability into the network, we will do it. The second is that it ultimately becomes a commercial issue. If we can start by selecting a smaller number of services for a smaller number of customers, that will have to make our job more simple and ensure that we do it better. This is one of the considerations that drove us to adopt our wholesale rather than a retail model.

COOK Report: Speaking of keeping things simple, did you say that your InterAgent network management software

can gather and keep statistics at the edges sending them to some central point only in times of troughs in the traffic?

Hanks: Yes I did. And the other piece that goes with this is that it's important for you to understand that services get delivered from a server on the edge not just in terms of raw bandwidth across the backbone. So the places from which we collect information give me information of a much finer grain. It is coupled something that is already doing computation and storage. Relatively speaking, it is more easy for me to accumulate this kind of information and transmit it back across our network during times of lower usage.

Now we still have some other issues that we're working relatively closely on with Cisco and with other organizations like CAIDA (the Co-operative Association for Internet Data Analysis.) In addition we are working closely with several other people in the industry on question how of you do real-time data monitoring in a way that makes sense. In the short run if you monitor the edges and the edges are based on servers rather than routers, you pretty much know what is going across your network because it is happening on a contract basis rather than on a bits-will-flow-where-they-may basis. Doing it this way makes it much more possible to build a network that will deliver the Quality of Services for which you have committed.

COOK Report: How would the question of using RED, which is a favorite of Sean Doran, fit in?

Hanks: It winds up being important. Cisco has a handful of technologies that are really important in terms of getting you the solutions for which are looking. Among them are RED, the CAR stuff and edge rate shaping, their Distributed Round Robin (DRR), their leaky token bucket and so on. All of these perform a particular type of function having to do with traffic policing.

If you ask why Sean likes read so much, you'll find if it is because he is operating in the Internet environment where you have fractal traffic that is filled with spikes and normally crest over a peak.

COOK Report: But you are a carrier with the of dark fiber capacity to sell, so how do you look at RED?

Hanks: As long as I have enough bandwidth to avoid congestion, I have no problem. But sooner or later I will run

out of bandwidth. It doesn't matter how much bandwidth you have now. As I said when you go back to look at the growth curves, you discover that the first derivative is positive and that the second derivative is also positive. That is to say bandwidth consumption is growing and the rate at which it is growing is also increasing. Now this starts to get rather scary when you think about the applications that have not been put on line yet. You have things like DSL that increase the end link capacity by three orders of magnitude. You know that the possibility of all this extra capacity that everyone is putting in the ground could be instantaneously absorbed is pretty high.

COOK Report: In other words, if we look at how this is he evolved over the last ten years, we can understand how impossible it is to predict what we will be doing application wise three to five years from now. What does seem to be a trend, is that somehow more and more of what was once done off the net seems able to be done more quickly and efficiently on the net.

Hanks: Yes. And to make that point, I have a T-shirt that I drag out from time to time. The T-shirt dates back to my graduate school days. It says "Third Annual Last Cable Pull Ever." Again and again you seem to pull so much cable you cannot possibly use all of it. But then you have a paradigm shift and, before you know it, the cable is gone.

For example look at the future where you have an average family of four that has for computers and two separately functioning HDTV's. Assuming that you only have one feed which is delivering everything, you could make a case for an OC-12 to that house sooner than you would think. Now take that hypothetical OC 12, multiplied by 150 million homes, and ask yourself with the backbone looks like?

COOK Report It looks awesome.

One of the things that stands out in talking to you is that your market orientation is primarily slanted towards audio and video — in other words towards isochronous traffic. Are you taking this approach because this traffic is the most difficult to handle? And because the real-time delivery of packets that do not have to arrive in precise order at the user's workstation somehow becomes easy by comparison?

Hanks: There are a couple of facets your question. One. it becomes very easy to

set-up two identical workstations on a LAN and have one of them play a Real video clip off a Real video server across the Internet and have the other play the same clip from our network and show customers a tremendous difference in the Quality of Service. It's the kind of approach that permits you to take an unsophisticated user and demonstrate that person that you are really doing something special.

The second thing is that in the technical community handling isochronous traffic is seen as being extremely difficult. Therefore if you solve that one problem to people's satisfaction, they begin to believe you when you say: by the way, we have solved these other three problems as well. So it's about buying credibility more than anything else.

Another thing is that we are motivated by looking at the functions that people would like to deploy today but that they can't afford to because of the facilities-oriented pricing model of the telcos. A lot of those kinds of services wind up having a strong video component. We believe that there are a lot of high value services out there for which people will actually pay a significant amount for on an as-used basis.

Multicast, QoS and Physical Coverage

COOK Report: What about multi-cast? How does it impact the complexity of growing bandwidth demand? If you can ever figure out how to price it and how to use it across network boundaries, perhaps it would serve as a tool to hold down demand?

Hanks: Those are the big issues. And they get down to being able to determine when an individual joins and then the leave us a multicast group. As well as whether they are authorized to do that and whether they have agreed to the build for doing it.

COOK Report: so deciding how to handle these issues, if you ran multicast on your network, would be as much of the problem for you with your overlay network as it would for single ISP?

Hanks: Yes. There are, however, some possible ways to solve this. There is emerging work going with the DMTF (Desktop Management Task Force) that is sponsored by Intel among others. Also there is this whole area called Directory Enabled Networks. Here the basic concept is that, as you authenticate yourself into

the network, you receive a kind of "capability pack" that describes what you can and can't do, your level of authorization and the amount of funds that you have agreed to expend.

COOK Report But in one sense this whole multicast thing is just another way of using bandwidth efficiently?

Hanks: Correct. Now to some extent Enron Communications obviates the need for multi cast with it's distributed architecture. The fan out is from the servers closest to the edge and the amount of traffic between those servers winds up being relatively small compared to the traffic that they are fanning out into general distribution. So multicast is both about efficiency and about the ability to join and leave a community of interest. The issues around this all get down to authentication and the directory enabled network stuff is probably going to take care of a great deal of that.

COOK Report By way of conclusion: what are the cities that you are either active in now or will be commercially active in within the next couple of months and where do you see Enron Communications headed over the next three to 12 months?

Hanks: Right now we have pops opened in Portland, San Jose, Los Angeles, Houston, Dallas, Chicago, New York City and Washington D.C. We have 3400 miles of fiber in the ground at about another 3000 under construction. Of course in addition to this we have capacity exchanges and fiber exchanges with companies like Frontier. Our ambition is to have a 20,000 route mile North American network up and running by the end of next year. That's 20,000 route miles of fiber that we either built bought or traded for and running through somewhere between 45 and 75 of the top cities in United States. We will also have an additional network that is extended on capacity swaps into other cities. At this moment I don't know what that will look like. The reason I don't know is because it will be demand driven.

Worldwide every time that Enron opens the ground we have an opportunity to evaluate the possibilities for building communications networks. Enron Corporation is in 167 countries around world and I think it is reasonable to assume that Enron Communications will be in a significant fraction of those within ten years. How that develops will be driven partly off of opportunity, partly off of demand and partly off of pure dumb luck.

Wither Telephony in IP Dominated World

IPtel Engineers Debate Merits of Innovation Versus Predictability in IP Telephony Protocol Design

Editor's Note: In late January on the IPtel Working group mail list, a question by a 3Com employee about the SIP protocol kicked off a discussion that migrated from the technical aspects at hand to a very informative discussion of the general role of the IETF process in Voice over IP protocol development. The debate pitted those who want the implementation of an entire suite of standards that would yield a total voice over IP solution, versus those who are happy to see more rapidly evolved modular solutions that may be packaged to serve varying needs and environments even if they are all less than 100% interoperable. The debate is also about accommodation versus innovation. Should we be striving to move the public switched telephone network to the Internet? Or should we use IP and computer technology to build an entirely new phone systems and set of services? The legacy phone companies and their mentor the ITU not surprisingly want a cautious conservative transition to globally interoperable standards at all levels. The engineers of the IETF favor innovation over predictability as they push for a process that can offer a mixture of computing and telephony services independent of POTs and the PSTN

On January 20 **Hanlin Fang** of 3Com asked on the IPtel list: My question is that: if you don't use proxy server in SIP, how can you get control and management capability comparing with H.323 gatekeeper system? If you use proxy server model, is there any significant difference between H.323 and SIP? Thanks for any comments!

On January 21 **Paul Jones** of the ITU responded: Hanlin, This is one of the very issues that I have struggled with. In the Internet environment, protocols consume bandwidth as they need. Latency can be high, but it is tolerable. The idea of introducing voice and video into the IP network means that there must be some control over the bandwidth, otherwise the quality of the call (if the call can be placed at all) will suffer.

Imagine this scenario: an earthquake strikes California. Everybody in the state

and out of the state attempts to place calls in and out of the earthquake zone to check on friends and family. Today's telephone equipment will allow a certain number of calls to succeed— all other attempts will fail due to busy circuits. With the use of SIP in an IP telephone device, there is nothing (that I can see) that would prevent everyone from attempting to place calls simultaneously. The end result is that all of the calls would suffer to the point that nobody could understand anything said on the call. Something must control the traffic. Surely the SIP designers have planned for this scenario. I would love to know how the situation would be handled in a SIP environment.

Voice and Video over IP Demand Control over Bandwidth

Jonathan Rosenberg (Lucent Technologies): There are several issues here that are being rolled into one. SIP does not address network QoS. That is left to other mechanisms, such as RSVP and/or diff-serv. With RSVP, routers would reject reservations once the capacity in a router is reached. With diff-serv, border routers will have policers to make sure that people abide by their SLA's; other mechanisms might be put in place to deal with overloads. The problem you mention is not SIP specific at all (which is why network QoS is not addressed in SIP). What happens if everyone tries to send email? Or enter chat rooms? Or make H.323 calls?

Given that, there are certainly mechanisms within SIP which allow you to restrict call volumes, independent of whether there is sufficient network capacity. A proxy can simply reject calls after the number of calls through it exceeds some maximum, or if the request rate exceeds some peak. Any criteria, at the discretion of the designer, can be used to reject a call. In SIP, the call is rejected by having the server send a 600 class response. Given that SIP INVITE's usually carry SDP (which contains media parameters, like codec types and possibly bandwidth usage), a server could also use this information to try to restrict "total bandwidth usage" in some abstract,

topology independent definition of the term.

What if there are no servers? What if users try and contact each other directly? In that case, you will need to make use of network QoS management techniques. This is also not SIP specific. If two H.323 terminals communicate directly, you'll have the same problem.

Henning Schulzrinne, (Columbia University and Lucent): SIP does not control network resource usage; that's the role of RSVP, diff-serv, etc. Indeed, it likely avoids the signaling overload problem encountered by traditional PSTN services as it may bypass all but the final end system. Since SIP can transmit limited out-of-band data, the California earthquake situation would actually be handled much more readily by SIP:

INVITE joe@sf.ca.us
From: mary@sj.ca.us
Subject: I'm fine but all the china is in pieces

[**Editor:** Earlier **Rosenberg** made the point that]: The problem mentioned by Paul Jones is not SIP specific at all (which is why network QoS is not addressed in SIP). What happens if everyone tries to send email? Or enter chat rooms? Or make H.323 calls?

Douglas Clowes: Well, H.323 bandwidth management in the H.323 gatekeeper will come into play. Of course, it will only limit H.323 calls and all the SIP, email and chat traffic will make it meaningless. I expect similar controls, introduced into SIP to have a similar lack of intended effect.

Rosenberg: Yup. Controls in a SIP server are good for localized policy decisions, (Joe can't make video calls), but for decisions based on global network state - well, that's really not feasible.

Clowes: Some sort of bandwidth management, and payload prioritization, is required to offer some level of control over QoS. That is, and ought to be, outside of SIP/H.323 and firmly within the underlying transport system which, at least

for H.323 (and I suspect for SIP), can be something other than IP. As I understand it, and Jonathan will correct me if I'm wrong, RSVP is supposed to "guarantee" bandwidth, while diffserv will prioritize the traffic. In the former case, the call should be good or fail while in the latter, if there is too much "high priority" traffic it will all be bad.

Rosenberg: Actually, its not if there's too much high priority, its if the amount of high priority traffic reaches link bandwidths. If you implement the EF PHB using a priority queue, and you expect on average 10% of the network bandwidth to be voice (which the ISP marks as EF), in an absolute catastrophe there is still 90% of the bandwidth left that the priority queue could usurp. If a priority queue isn't used it might be less. So, in a network with lots more data traffic than voice, things might work quite nicely - except for the poor folks trying to send email :)

Editor: To **Rosenberg's** statement: What happens if everyone tries to send email? Or enter chat rooms? Or make H.323 calls?

Jim Toga (Intel) responded: Call admission by H.323 Gatekeepers can be used in this case to match provisioning with capacity. The behavior might exactly mimic the current PSTN.

Scott Petrack: There is unfortunately a fundamental difference between the PSTN behavior and what an H.323 network can do. When the PSTN decides that there isn't enough capacity, the end user gets a Network Busy signal and there is nothing else s/he can do to affect the telephone network. This is because of the fundamental distinction made in the PSTN between "NNI" (network-network interfaces) and "UNI" (user-network interfaces). The signalling which tells the switch "No you may NOT make that call" stays entirely inside the network, and the business contracts and regulation that bind operators guarantee that when that happens, the end user will not get anything but a busy signal. He will not be able to mess things up for anyone else.

But H.323 signalling is not NNI — it goes all the way out the endpoint. So if a gatekeeper tells an endpoint (which might be a terminal) "No you may NOT make that call", and the endpoint can get via other means (say email) the RTP address of the remote endpoint. A user at the endpoint can blast away. It may sound bad, and not many packets might get through, but he can certainly inject the packets into

the network and still mess things up for everybody else.

There is no way within H.323 to connect the permission that a gatekeeper might or might not give, with the opening and closing of IP connectivity to the endpoint, and there is in the PSTN.

Note that, even if such a connection existed, it would not really help, because there are many different kinds of traffic which affect capacity, not just H.323. Even if the Gatekeeper really could stop me from sending RTP traffic from being sent when it forbids an H.323 call to complete, it won't be able to stop me from sending Email, ftp, etc. The bottom line is that Gatekeepers are not enough, unfortunately, although the analogous mechanism is enough in the PSTN — because the business contracts and laws are there and because there is only one single application running (voice).

Rosenberg: SIP servers can also do "admission control", by rejecting requests (see my previous message). So, whatever logic you have behind the admission decisions in RAS, you could put in a SIP server too. However, its not clear how either a SIP server or H.323 gatekeeper would be able to actually make a decision which mimics the PSTN. In the PSTN, the circuits flow through the switch which makes the admission decision, and the switch sees every setup message for a circuit that goes through it. In the Internet, media doesn't flow through gatekeepers or SIP servers, and neither will have knowledge of network topologies or other applications. So, if there is lots of web traffic on an intranet, I don't really see how a single, application specific box could know how much bandwidth is actually available on all links between a caller and callee. Only routers can do that, and that's why QoS control is done with things like RSVP and diffserv.

Complete Telephony Solution Using SIP?

Editor's Note: From here the discussion broadened out and went in the direction of what type of philosophy should be brought to protocol design in IP Telephony.

Jones: It is certainly acceptable to say that SIP should not manage the network resources given that its function is a "Session Initiation Protocol". My question really was directed more at the overall solution, which involves SIP, than SIP by

itself— and you have answered that question. However, there are other, similar questions that need addressing as well.

Are there any efforts going on inside the IETF or in other groups to define a complete IP telephony solution using SIP? That is, are there any documents in existence or under construction that provide guidance to implementers or service providers that address bandwidth issues, user location, security, roaming, etc. Is anyone addressing the issues of lawful interception or access to emergency services in the SIP environment?

Rosenberg: Some of these issues are being addressed; SIP itself addresses security, relying in part on existing mechanisms, such as TLS or IPSEC. User location is one of the features of SIP; work is going on in the enum group (not yet a working group) around doing it for telephone numbers. I know of efforts regarding roaming services and emergency services (a REALLY hard problem in an Internet context), but nothing at the standards level.

However, there is currently no IETF effort to specify a "complete solution" for IP telephony. In general, this is not something IETF does. It is good at making nice, general purpose, modular protocols, and the job of putting it all together for a complete solution is something that the service providers and vendors do, as there is more than one right way, and it leaves room for growth and value add.

Clowes: And there are any number of almost, but not quite, entirely non-interoperable implementations. While there is more than one way that might be right, there are far more that are not right. Product-differentiation, market-fragmentation, or customer-lockin: call it what you will.

Henning Schulzrinne: There may well be value in agreements beyond protocol specs. However, there's often the suspicion that many of these industry bodies are clean-up crews for standards that could have been clearer, been based on more implementation experience or more layered/modular to begin with. The abundance of these organizations also raises the cost of entry, so they are probably at best a necessary evil, not a model to strive for. Note also that the traditional telecom world is hardly a model to emulate here, if you recall trivial things like the number of different phone plugs, dialing plans or ringing patterns (see a previous discussion on rem-conf) in the world or the

number of ISUP/TUP/other signaling variants. At least email and web access from India to the US works without gateways and translators and interoperability agreements and I don't have to be certified as compliant in each of those countries.

Rosenberg: [The multi-protocol web] seems to have worked so far; there are many applications on the net which involve a mix of protocols. The web is certainly one example - it relies on http, of course, but also on back end directory services for dynamic content, HTML, Java, Javascript, HTTP-CGI, cache replication protocols, and lots of media types and plugins. There is no one spec which tells vendors what to implement in their servers.

What Should the Goal of IP Telephony Standards Design Be?

Jones (ITU): In order to deploy a global IP telephony solution based on SIP and related technologies, we must have a standard that addresses the "loose ends" that the protocols by themselves do not.

Clowes: And there are so many ways that these things can and do break. With the latest browser version available, I still got a site that told me my browser was too old and I should get an update before I come back. It breaks.

Rosenberg: I'm not yet convinced that we "MUST" have such a standard. I agree there must be solutions to the various component problems, but not all in one, birthday wrapped solution. Grandiose, all encompassing solutions have a bad habit of never making it.

Clowes: The solution seems to be the "Implementation Agreement" where, instead of the standards body picking and choosing, a group of vendors, implementors, or customers (or all three) get together and make the choices. Sometimes each group gets together separately and we get three or more IAs. IMTC has an IA on VoIP telephony. ETSI TIPHON is a selection and tightening of the standards. There's also iNOW! and TIPIA. And some others that escape me.

Henry Sinnreich (MCI): We exchange here e-mail and web pages by using IETF standard protocols. Am not aware that our applications, whatever they are, went through implementor agreements and the like.

Paul Jones (ITU): But even mail and web services are not without problems. Take a web browser written in 1995 and visit some of today's more elaborate web sites. Create an e-mail with Microsoft's Outlook Express using HTML and send it to a user with a mail client written in 1995 (heck, send it to a Netscape user with up-to-date software).

It is true that most of today's mail clients can understand SMTP, POP3, IMAP, HTML, MIME, etc. It is also true that most of today's web browsers can support HTML 4.0, CSS, GIF, JPG, PNG, WAV, AU, MPEG, RealAudio, etc. However, it has not been without pain on the users—they have had to upgrade their software every time a new "whiz-bang" feature was added. Be prepared to upgrade your web browsers again soon as XML and XSL displace much of HTML and CSS and MP3 becomes a new important audio format on the web. We cannot and should not expect our telephone system to act this way.

Scott Petrack (formerly of VocalTech now of Metatel): With all due respect, this is really just because we are in the absolute infancy of "Internet Based communication". The Web is less than 6 years old, for goodness sake, and the PSTN is well over 100 years old. I believe that there was quite a long time in the early days of the telephone when you needed to have 2 or more phones on your desk if you wanted to talk to people on different networks.

Paul Jones: People will be investing billions of dollars into IP telephony solutions—they had better talk to each other. We must agree on how to authorize users when delivering calls to gateways. We must agree on how to locate users or resolve addresses. We must agree on QoS issues (use RSVP?), billing, roaming, access to emergency services, etc., etc.

When I go to the store to buy my IP phone, bring it home, and plug it into my hub, it should allow me to dial my grandmother's phone number and her phone should ring. The mechanics of accomplishing that are complex, and we must have agreements on how that should be done at every step of the way. In addition, I expect that my IP telephone will continue to work for more than four years.

Creating an "Implementation Agreement" as Douglas suggests is not a bad idea. It was not necessary with other protocols because they were simple, the implementations were software-based and free (so cost was no issue), and developed

without interest from parties investing billions of dollars. The last point is the key—when people spend billions of dollars, they expect it to work. Specifying a bunch of loosely related protocols and letting people build several implementations, picking and choosing incompatible components to address requirements, and waiting to see which one is "coolest" and declaring it the winning combination will not work in the IP telephony space.

Tim Dorcey (IXL): I dunno, I think it could work out pretty well that way. If my IP phone works just like my current telephone, then I don't want one. I can't speak for the people spending the billions of dollars, but I am most interested in IP telephony for the innovations it will bring. And, I don't think anyone would suggest that standards committees are fertile ground for innovation; it's just not their mission. Also, let's not forget that it took many years for the existing telephone system to reach it's current level of maturity.

Sinnreich: Agree, it is not imitating the telephone network, but the platform for innovations that makes sense. Also, then it will not be IP telephony any more, but rather a rich mix of services, of which voice is just another component IMHO.

Clowes: Henry, The number of people that I hear cursing Microsoft, and Netscape too, because they don't implement each other's features *today* is not small. Each browser only gets best effect when used with their own server. What works in a Netscape browser, only when served by a Netscape server (and v.v.) is an implementation agreement within the vendors browser and server teams - same for MS.

Often, when I send a message to a mail exploder, I get back one or more non-delivery notices. Sometimes I get one mail message more than once. I suspect I get some messages not at all. On the latest available version of MS Explorer, I loaded a web page last week that told me my browser was too old. It was wrong. It was broken, just like the mail system.

I can't send SMTP mail messages when I'm "off-net" because my provider will not allow it, because it has inadequate security. I can't use NTP/SNTP through firewalls because people block it. There are no end of ways the Internet, and IETF protocols, break or are broken.

It really depends on your expectations, and those of your paying customers, whoever

they may be. From consumers who expect to pick up a handset, push a button and talk to mom, without restarting their handset and redialling three times. From service providers who want to be able to buy a new box, and just plug it in, without hiring a consultant to tell them what support servers they will need and will it be compatible with their existing infrastructure.

Incidentally, e-mail and web exchanges only occur under two specific circumstances: either you are lucky that there is a chain of commonality, or your software has implemented a great number of almost, but not quite the same protocols. My old POP/UUENCODE mail client has problems with the IMAP post office, and the MIME attachments. The free mail clients that support them are brain dead, so one day I'll fork out for a new one. Until then, I'll plod on and curse these IETF protocols. [The nice thing about standards is that there are so many of them from which to choose.]

Rosenberg: From the box builder's point of view, he wants to build boxes with any damn standard he want to, and update to the latest protocols to remain "whiz bang". From the box buyers point of view, he wants to buy boxes from different vendors and have them work in his existing installation. He doesn't want to find that of the two new boxes he bought, one uses RADIUS while the other one, with the higher serial number uses DIAMETER and that neither works with the same models he bought last year from the same vendor that use TACACS+.

Ping Pan: Sure. That's why some of us (from multiple vendors) are working hard toward a single standard protocol that can provide accounting, authentication and admission control for many emerging internet applications (IP telephony included). Actually, some of us are not interested in "whiz bang", rather the market itself.

Clowes: Some people want to sell boxes. Some people want to buy systems.

Ping Pan: Yep, we want to sell the boxes now so that the customers can get into the service games as soon as possible. As the protocol and development evolving, the sellers can always send software versions to keep the customers up to date. SIP and QoS

Rosenberg: Question: besides the specification, what prevents a client from ignoring them anyway? A big theme on

another thread has been that clients often disobey or mis-implement portions of a spec. IP is IP, and two hosts with globally routable IP addresses can always communicate. Seems like a bullet item for a feature in fact: no call rejections - the user's software just tries to call the other user directly, ignoring the admission rejections.

Clowes: Of course, if it doesn't implement the spec, it doesn't conform to the spec. People could ignore the standards and implement their own protocol. I think that's happened before. Once or twice. Let's see now, a firewall with proxies that only let through known protocols and check with the gatekeeper for H.323 admission. Scarey!

Francois Menard: No no, The whole H.323 hypothesis is that you can control the end-points and force them to use a given gatekeeper. This is as long as the end-users are not capable of figuring out how to change the config of the end-points to bypass the GK...

I've always maintained that on the Internet, it is not because you route a call that you can bill for it. In H.323 on the Internet, you can bill for: 1. QOS for all applications at the same time on the host 2. PSTN access through a metered H.323 PSTN gateway 3. Apps on GKs. (i.e. there is some value in a GK translating a phone number to a dynamic IP address)

So I do not believe that basic H.RAS stuff actually makes it possible to derive a bill. GK-routed call models are a billing mechanism, not a bandwidth preserving thing. So, by definition, it is only possible as long as you can force the end-points to collaborate. Same for a SIP server-routed call models too. But then, SIP makes no assumption about the rationality of SIP-server call routed models to create illusions about the real problems of IP telephony.

Scot Petrack (referring to statement of Douglas Clowes: In H.323 you would probably use "gatekeeper routed calls" - that is, the gatekeeper itself is the signaling proxy that only lets though known and permitted signaling, and tells the relevant RTP firewall what to let through. (But this is just quibbling with the words you use).

And if the signaling is SIP then you replace "gatekeeper" with "SIP proxy server." Although you can do it this way, unfortunately you have created — the phone system! That is, the RTP firewall looks a lot like a switch, because you have

to make certain that all the RTP traffic goes through that box. The originating endpoint will never see the RTP address of the remote endpoint, but only the address of the RTP firewall. So all of end-to-end routing has been thrown away — just like in the phone system.

Of course, similarly, we will need, SMTP firewalls and special DNS servers which can give out addresses of MX machines which can do QoS email, and other special DNS servers and HTTP firewalls which can do QoS web-browsing, etc. etc. In the end we will have overlaid 1000 different parallel Internets on top of the current Internet. I know many people who share this "dream".

A better, more scaleable, goal, is to separate out the QoS infrastructure from the needs of any particular application, and instead of having H.323 Gatekeepers, HTTP Gatekeepers, SMTP gatekeepers, etc. etc., just having an infrastructure for QoS that allows applications to seize the required Internet resources without any reference to any particular application.

It may end up being impossible to do, but it's the only hope that we'll get QoS for all the Internet applications that need it.

It may well be that there is still a need for the single-application servers, in order to enable really dumb hard-coded appliances that will not want to deal with the QoS signalling themselves. But these servers will need to hand-off decision-making to the real QoS infrastructure.

Summing Up: Flexibility vs. Degree of Effort

On January 25 **Jay Batson** (President Pingtel Corp.) wrote: This thread is effectively pointing out the fact that there are philosophical differences between people. Not that one is "right" and one is "wrong." And that the protocol must allow accommodation.

The issue is flexibility vs. degree of effort. When everyone is forced to agree and do precisely the same thing, progress speed slows from the speed of innovation to the speed of agreement. I, for one, *prefer* to live with some imperfection in exchange for faster innovation and change. I prefer Internet-speed to Telecom-speed.

The very fact that two end nodes can agree to do whatever they want across an IP network, and that the IP network provides

simple transport (hopefully soon with a QoS component), permits an incredible rate of technology advancement. Last Thursday I learned about ImagineRadio — a ‘net-based “radio station” that I can customize to my tastes, and listen to with a Real Audio player. Supported by text ads (no voice-ads.) Some guy thought of it, and I simply added RealAudio to my web browser to support him. I didn’t need UUNET to “agree on a implementation of personalized radio station traffic carriage.” And no, it isn’t as good as FM quality sound... yet. But I listen to it anyway because the music choice fits me better than anything I can get over the airwaves. And it all happened in a few months simply because the basic protocols existed — not a standardization of how the protocols are supposed to be used. Where’s PSTN “radio?” Non-existent.

Today, I can choose to deploy to my company an email application that does *straight* text in a predictable, reliable, manageable way. Or I can choose to deploy Netscape/Outlook email and have graphics/rich-text displayed in-line, and pay for a higher degree of management/software cost in the process. Or get some other feature that some vendor provides, simply because POP and MIME are so elegantly extensible — X-mynewfeature: application/mynew-whizbang-thing. We here at Pingtel have chose the innovative email, despite its warts. Of course, in either case, email gets through....

In contrast, look at how telco/cable-driven “interactive network” experiments have gone. Closed environments in limited regions with bogus, thinly disguised restaurant advertisements substituting for “entertainment” guides. Ask any business manager who has been in the “Smart Phone” business about telco-driven “interactive phone” trials where the telco provided the “content,” email servers, (Minitel notwithstanding.)

Compare that to boston.sidewalk.com, or www.boston.com, where there is competition, incredible innovation on how to organize information. This all works because of the genius of HTML/HTTP/Javascript/... on an “open,” public IP network. Extensible, simple, text-based protocols that anybody can innovate around. And (let me cast a barb here) do without buying an expensive ASN.1 compiler... ;-)

Paul may rightly feel that he wants a higher degree of stability and predictability. That the protocols should provide the same experience as the PSTN. This is his choice.

He should find service providers who *focus* on this, and *eliminate* products from their network (and those to whom they connect) that would replicate the PSTN experience using IP technology. His service providers should *prohibit* technology that contains the degree of freedom that I would choose.

However, I should be able to choose from service providers and products that are to IP Telephony what UUNET and Netscape are to the Internet: they provide reliable, basic, and extensible infrastructure that can be extended at will between consenting parties. I *want* things to work the way HTTP, HTML, SMTP/POP, ..., all work today, warts and all, because I want the innovation.

I *want* an ITSP that will, at some point in the future, provide me with the ability to buy/install (at my site) my own leading-edge application that will take in my SMTP mail, and use “3rd party signalling” with my ITSP (via SIP/MGCP/whatever) to set up a voice phone call between me on my IP wireless (premises) phone (from Symbol Technologies) and the sender of a SMTP message, via a service-provider-run IP/PSTN gateway, all because the words “CALL ME” were in the subject line of the SMTP message.

And I want the protocol open enough so that Joe’s Software in Palo Alto can write the application, sell it to me, and have it work with my ITSP’s (probably Henry’s, er, MCI/Worldcom’s ;-)) IP/PSTN voice network. And if it doesn’t work occasionally because it’s still immature, I’ll put up with it because the infrastructure — read “the protocols” — permitted the extensibility *without* requiring absolute, rigid adherence to the *way* in which they should be used.

In the End Once Again "Bell Head" vs. "Nethead"

Paul Jones says: We cannot and should not expect our telephone system to act this way.

Batson: There is another “we” who feels an entirely contrary way: We cannot and should not constrain the IP voice systems of the future — they can, and *should* work the way web browsers/email/... all work today, with the need to upgrade/install software to get the new whiz-bang features. The price of innovation is that you have to work to keep up. Give me the work anyday as long as you give me the innovation.

Give me choices in how to locate users. Give me choices in how to select IP/PSTN gateways. Let me choose which QoS choice I want (by giving ITSPs choices in products to deploy.)

I want a *choice* in *what happens* when I pick up the phone to call grandma. I *do* want it to connect. I do want some base level functionality — just as HTML/SMTP/... all do their basic job on the ‘net. But let extensibility roam free. If in the process, products don’t work with each other, fix ‘em. Everything does *not* need to be perfect day-one. This has plenty of precedent in the IP world. Back in the early days of IP dialup, when we called it “Remote Access,” various dialup servers didn’t all work with the same client software. We all used PPP, but “differently.” The market didn’t buy it. Customers demanded that a Mac PPP, a Win 3.1 PPP, and a Win95 PPP all be able to dial into the same server. So we created the “PPP Forum,” and did PPP “bakeoffs.” (Just like the VoIP Forum....) Today, do you have any *clue* whether you’re dialing into an Ascend box or a 3Com/USR box at your ISP? No.

This is not intended as derogatory, but to my mind, this is a pure “bell-head vs. net-head” discussion. The net-head approach to protocols has certainly created some great new innovations in the last 5-10 years. I’d sure hate to take that away from IP telephony’s future.

Editor’s Note: Finally on 2/23/99 there was this announcement on the IPtel mail list from Mo Zonoun of Nortel. ANSI TIA TR-41.3.4 has recently started a new project “Performance and Interoperability Requirements for Voice-over-IP Telephone Terminals”, to produce a standard for design of IP phones. This standard is to include all aspects of such designs, from physical connectors, to speech coding, audio quality and echo management, to control protocols. Initial focus has been on use of H.323 and Megaco. The kickoff meeting was held on 15th Feb 1999 in Sarasota FL, with 21 attendees including Lucent, Microsoft, Nortel Networks, Siemens, Cisco Dialogic, NEC, Bell Canada, and a number of small companies. The next official TIA 41 meeting is slated for the week of May 17th, with the TR-41.3.4 working group meeting to finalize as much of the spec as possible. An interim face-to-face meeting of the TR-41.3.4 working group was also proposed, tentatively mid-April. The spec is scheduled for completion by Sept 1999, using a TIA fast track process. For more info go to www.tiaonline.org

Business Model of a Third World ISP

Dileep Agrawal Explains the Dynamics of Building World Link, Nepal's Largest ISP

Editor's Note: Dileep Agrawal runs World Link, the largest ISP in Nepal. We interviewed him in Kathmandu on November 9, 1998.

COOK Report: How did you get into the ISP business?

Dileep Agrawal: I was born in Nepal but I spent much time in the US because my father did his PhD in biochemistry there. He also taught in the US. When we came back to Nepal, I stayed here for eight years to finish high school. At that point I returned to the US for college. As far as my education is concerned, I was in many places. I started off at a college in New England which turned out to be too expensive. So I went to a small college in Montana, Western Montana College, in Dillon.

COOK Report: Then you must know Big Sky Telegraph and Frank Odasz?

Agrawal: yes, I was involved with them. Anyway, I attended Western Montana College for one semester but I didn't really like their academics. Fortunately, I got a scholarship from Bates College in Maine to study there. So, I transferred there and got a degree in biochemistry in 1996. While at Bates, I took off for a semester and went to American University in Washington, DC for one semester of some business courses that I wanted to take. During the Summer of my junior year in college, I came to Nepal on vacation. I took a look at the Internet situation here and how things were being done.

At that time there was one service provider who was charging outrageous rates. For example, did you know that you're charged by kilobytes, the size of the e-mail. They charge about 2 rupees right now. However, when I came here, they charged 50 rupees (about \$1) for 1 kilobyte. I thought, Wow! This is a very profitable business to get into. Why should I think about finding a job in the US after I finish college when I can probably get this business started? So I collaborated with my brother who was an electrical engineer in the US. We discovered the technology; how it's done and everything else in 1995. In other words, I was aware of the technology from my experience in college in the US, but it was when I returned to Nepal that I first saw the opportunity to become an Internet service provider.

During the summer of 1995, my brother and I set up a small server running BBS software which allowed us to give UUCP

accounts to subscribers. We started off by charging ten rupees a kilobyte which is still a lot. But we slashed prices down and attracted a lot of small users, and ex-patriots.

Early UUCP Connectivity in 1995

We were able to give very good support. For a UUCP-based store-and-forward email system, there's not a lot of infrastructure requirements. It's not that complicated. Therefore it wasn't hard to set it up during the summer. In January 1996 I returned to college in the US and my cousins took care of the business in Nepal.

My brother had a computer in his home in Seattle. He would dial into our computer in Nepal rather than us dialing directly into the service provider. We would then exchange emails. Then he would exchange files with the ISP. The advantage to this set-up was that we could compress the e-mails (zip them up) and use non-standard modems. Service providers usually have Hayes-compatible modems but I discovered during my summer research that for noisy lines there's a company that had a proprietary protocol called Turbo PPP. This protocol works much better than the Hayes protocol over noisy international lines. This is how we operated. My brother had one of these modems on his end and there was another on Nepal end. We compressed the files and sent them to Seattle. He'd unzip them and send the emails to his local ISP.

In January our competitor got a leased line to Singapore and was able to provide full Internet connections. But they had a lot of problems with it because it was only 19.2 Kb/s and they weren't able to provide all Internet services. And their prices were still outrageous. Not even the ex-pats [ex-patriots] who usually had money in their pockets could think about getting an account. Only the rich INGOs (International Non-Governmental Organizations like the UN) which had projects to support could subscribe to it. One of the INGOs was paying our competitor \$2,500 a month for a leased connection. There were people who were desperate for connectivity.

Our business kept flourishing. People were very happy with us because we were able to provide personalized, door-to-door service. We ourselves would visit the customers to install the software and train

people, or we would send someone to do it. There are ex-pats in the foreign community living in Nepal and they really like that. Our customer base grew by word of mouth.

So I finished my senior year at Bates and then returned to Nepal in June 1996 and made some improvements to our server. Our business was going fine and we were satisfied with it but we realized that if we wanted to remain competitive, we'd have to get a leased line. We'd have to be online because at that time our e-mails reached their destination in one or two hours, namely whenever we made the call to the US. But our competitors would send e-mails out immediately. My cousin Shyam got a scholarship to attend INET95 in Montreal. I had planned to go there straight from Bates but I thought that Shyam would get better exposure if he went. Shyam met people there from TeleGlobe International and that helped us a lot. We kept in touch with each other.

COOK Report: Was Bob Collett with them in the summer of 1996?

Agrawal: He's been there ever since they started their international program. But the TeleGlobe people were very aggressive. We only gave them about \$4,000 worth of business a month but they looked at us as a gateway into the region. So they were very eager to get us to sign up with them when we finally had the cash flow to support it. Remember, we had to pay \$4,000 to TeleGlobe and another \$8,000 to the local PTT in Nepal for providing half of the connection. This was in 1996 but the rates are about the same today. We know the rates in Nepal won't go down because our PTT has a monopoly [in the local loop and international connectivity]. And we haven't pressured TeleGlobe because they're already providing us with a very reasonable rate, i.e. \$4,000 for 64 Kb/s using satellite connectivity and their own fiber connectivity that they have in Europe. This is a small component of the total cost. We're happy with them. And remember, they connect and terminate us more effectively to the rest of the world.

TeleGlobe even gave us the facility under terms by which we wouldn't have to pay anything for the first three months and we could pay quarterly after that. This would allow us to build up a cash flow. We were a very small company then with only five employees (including Shyam and I) and the idea of spending \$12,000 a month is mind-boggling-even for a company in the US. But we were in Nepal where such an idea

was unthinkable! When we told people we were spending this kind of money, they just couldn't believe it. But I took this bold step anyway.

In the meantime, I spent three months (November 1995 through January 1996) studying Linux, getting on discussion groups like INET Access and bugging people for help. Then we finally decided we had to go to a full internet connection. Before this we had been running a DOS BBS. But we figured that would not work. We had to set up under Linux as an operating system. [Windows] NT was too expensive and required too many resources. So I took it upon myself to learn a completely different operating system, setting it up for an ISP, and getting discussion groups to help me. It worked very well. By February 1997 we had everything working, including our server, our interconnection and our routers.

From that time on things got really wild. They went too fast for me to understand. When we came into the market we slashed prices. However, this backfired on us because our competitor slashed his prices even further. We ended up matching these prices later on. But the end result of all of this was that more people got involved in getting an Internet account because it became affordable. This brought in a lot of business for us. Now we are at the end of 1998 and find that we have to get more capacity because a single 64 Kb/s leased line isn't enough. We're doing a lot of heavy caching. What we really need is a 512 which we're looking at now.

Internet Availability Brings on New Classes of Users

COOK Report: You said you reduced prices to bring in more customers and increase the size of your market. Who were some of the people who came on board? The tourist business obviously. Did the CyberCafe start then? Give me some details about the business components you developed.

Agrawal: Previously our customers were primarily ex-patriots, that is Nepalese who lived abroad and wanted to stay in touch with people back home in Nepal. Soon the foreigners living in Nepal became another component of our business. They were the only private home users using the Internet at that time. What's happened since then is that people have become more aware of this technology. For example, you may have heard of the onion crisis here which happened when we had a very bad onion crop and production almost stopped. Onions are a big part of cooking. So importers got on the Internet and started looking for onion exporters in various countries. Businessmen who have learned

how to use this technology are definitely benefiting from it. We can see it. Those people are logged in for HOURS. And they wouldn't be spending the money unless it benefited them in some way.

Watching the price of gold is another thing businessmen are interested in. What's the current price of bullion in the international market? They can go to the London metal exchanges, check the rates there and adjust our local rates accordingly. People use it a lot for this as well as finding out what's the current spot rate for dollars in any part of the world.

After the business people discovered this technology and began using it, there were people doing research in the university who began to use it, too. This is the segment that's coming up now. But the tourist trade is also important because it's the backbone of our business.

COOK Report: What about Nepalese business people who run hotels, promote tours and/or trekking expeditions, etc.? When did they come on board? Before or after or at the same time as the other business people?

Agrawal: They came before but they weren't really the first business people to sign up. (The ex-patriots were the first.) They came on when they realized that their clients-clients like you- found it easier both to find them on the Internet and to communicate with them via e-mail. That's when they realized that they should get an Internet connection or put up a web site. Before that, they didn't really know or care about it. Almost every big hotel in Nepal has a web presence. They have some provision that offers a discount if you book over the Internet and they accept what comes over the Internet. They realize the Internet has potential and they need to be connected in order to tap into this market. We do have a web design business but we do it on the side and haven't really concentrated on it. But we have it if clients request it. We use our own staff or hire the expertise to do it.

COOK Report: Suppose an importer wants to build a web site. What are his choices? He can build it himself or come to you. Are there also university courses and students who can do it?

Agrawal: There aren't any university courses but there are private institutions that have courses. But these people usually associate the web content with the Internet so they call us or the other ISP for help. So when we get a client we usually get their web business as well. They just find it easier to deal with us. For instance, someone came to us the other day and said, I want a web site for my company that looks just like the BMW site. So we'll do it or hire the expertise in the market. We have three full time employees who can

work on web design and if we have surplus work, we outsource it. If the demand for web design grows, we can always hire a fourth person. We'll have to see which segment of our business is most profitable. Right now we just haven't been able to concentrate on web design. We have to take care of our main business more. If we get a job that's really big, we just outsource it to other people we know can do the work.

Infrastructure Issues

COOK Report: You said that all of your constituents realize that the Internet is fairly expensive compared to what they're used to, but they do find that it benefits them. Where does it go from here? Are all of your customers dial-in? Given the \$8000 a month local loop cost is there any other practical way to do it at this point?

Agrawal: I think that you misunderstood what I said about the local loop. Our circuit comes from Montreal to Denmark. From Denmark, TeleDenmark uplinks it to an Intel satellite which lands in Kathmandu in earth station owned by Nepal Telecom Corporation (NTC), the local PTT. The cost is shared in the following manner. Since it's shared circuit, 64Kb/s out-going from Kathmandu to Denmark is provided by NTC and 64Kb/s in-coming from Denmark to Kathmandu is provided by TeleGlobe, a half circuit. Two circuits of fiber optics from Montreal to Denmark are provided by TeleGlobe. Plus the 64Kb/s Internet connection cost in Montreal is provided by TeleGlobe. SO NTC isn't charging for the local loop. They're charging for the segment they provide on Intelsat plus the earth station facilities they provide for reception of the signals. In short, it's the traditional half circuit, in that the American half of the half-circuit is much cheaper than the foreign half-no matter what kind of circuit you're talking about.

But the local loop itself is very cheap in Nepal. It cost 10,000 rupees (\$150) for one year. What they give you is a pair of copper wires and you have to put in these HDSL modems that can pump 2 Mb/s through it, or 64 Kb/s, or whatever you want. When we began offering our service, we had to put our own modems in. We started with modems by Rad. They're an Israeli company. Many companies make these modems. If you have a pair of copper wires and you want to send 2Mb/s over them, you just put these modems on the copper wire. However, there's a distance limitation.

COOK Report: How close are you to the earth station?

Agrawal: The circuit from the earth station comes through the local exchange which is the building next door, and it comes through their existing digital circuits. From

there, it goes through the local loop to our office which is about 3-4 kilometers from here. So from the local central office to our office is HDSL. Our servers are on the other end of that.

We encourage people to take leased lines. We have about five leased line customers, like the British and Netherlands Consulate and a few others. We give them the same bandwidth they would get if they dialed into us. They save on the local loop cost. For example, in Nepal you pay one rupee per minute which is what the PTT charges. So if you're logging in for 3-4 hours a day, it would be cheaper to get a leased line and hook up with us.

While there are customers getting leased lines, most of our customers are dial-up. Even so they don't use it a lot because it's still expensive. They use it for what they need to get on and get off. If you're a professional user, you probably check your email and get off. For the common people it's really not affordable.

We have four more offices around town in addition to this one. They're primarily for money collection. In the US, people are used to a pre-pay system for the Internet. In Nepal, we spoil our customers by billing them. But then we have to get people to pay. Since there's no way to control it, we have to knock on their doors and ask for the money. It's difficult. In fact, we're thinking of doing away with this system and starting a pre-pay one. People will complain but it's a business decision. We'll have to do something like that because it's getting more and more difficult as our client base increases.

I can't believe it but we have about 50 people working here. Ten of those people are involved with accounting, payments and customer follow-up. Two or three employees take payments from walk-in customers. There are about 15 people roaming in the field at all times taking care of customers, like checking on the leased line connection or handling complaints. About seven people provide tech support on the phone from 7:00am to 8:00pm but the time isn't really fixed. Sometimes someone is here until around 9:00pm. We're closed on Saturdays.

We have another service called fax-to-fax. You can send a fax from here to almost any country in the world for 20 cents using the Internet. It goes through the Internet to the US and then gets faxed from there. We have a separate line for that and we help people to use it in the field or to track faxes on the Internet.

COOK Report: Is the Cyber Cafe by the bus stop the only one in town?

Agrawal: Yes. That's the only one we have. It's been there for about 6-7 months. We opened it primarily to have a place in

the tourist district for contacting the tourists who go there. It does a good business because it's centrally located and people need to use the Internet. There are three ISPs here in Kathmandu. After we got started, one more ISP came in last year. But they're not really active in their business. They have other businesses to take care of. For example, they first started the fax-to-fax service and that gave them a boost because it was very cheap to send faxes for so little money compared to what the PTT charged. That really set them up and consolidated their business. So they've been concentrating on fax-to-fax services. Providing Internet connections is just a side thing they do when people ask for it.

COOK Report: I get the impression that you built up your business from your original cash flow. Is that true?

Agrawal: That's correct. When I came here, I brought one computer and a few modems. That was it. That's how we started. We didn't have to inject any capital after that or worry about bank loans.

COOK Report: That's how it used to be done in the US. But now the field is so populated with the big guys that individuals don't have a chance.

Agrawal: My advice for any of the entrepreneurs in this region is, if you have a chance, study what's happening in the US. If you adopt and implement it here, then you will do well. Whatever happens in the US will happen here 5-6 years later. So if you learn from the experience in the US and implement it in Nepal, you'll definitely succeed. I saw the business model that start-up ISPs were using in the US and I tried to implement it here.

Of course, that business model is no longer applicable in the US because you have so many bigger players, and, if you don't have a great deal of cash at the beginning, you can't succeed anymore. But by taking advantage of the window of opportunity at the right time, we've been able to position ourselves such that, barring any disaster, we can continue to play a major role here. Timing was everything for us. If I hadn't started in my junior year, we probably wouldn't have had the opportunity to grow. What happened here is that, because it's a small market, once you start up, other people who had been thinking about it just sit back and won't come into the business. Why? Because of the limited depth of our economy, the market is small and everyone has to share the same pie. And the costs are substantial, namely \$12,000 a month fixed cost just for the circuit. That's a huge amount.

COOK Report: India prohibits any Internet telephony. What is the position of Nepal Telecom Corporation on this issue?

Getting Along with the PTT

Agrawal: Let me answer that in the context of talking about the government issues. That's very important for these kinds of things. At the time we started, there were no regulations-none whatsoever-because the government didn't have a clue as to what was happening. It's worth mentioning that our predecessor Mercantile Office Systems was very brave in that they took a risk to start this. They paved the way for people like me to come in. I don't know if I would have taken the first step on my own. But these guys had already taken the risk by starting first. And it appeared there was indeed a market for the service and that the government wasn't going to do anything about it. We could see that they weren't going to come and shut us down. So we went ahead and started our service.

I don't remember the exact date but it was last year, around October of 1997, that the government realized what was happening. There's an annual trade show on information technology in Kathmandu in February. That's when people get to show off whatever they have. At this trade show both we and our competitor were able to show people what the Internet was. It was at this point that the government realized that this thing [called the Internet] existed and they had to do something about it. That's when the ministry started to formulate policies. In September they approached us and said we had to apply for a license. We answered: OK what are the guidelines?

We had to explain our set-up to them and tell them about our capital and things like that. So the government imposed a 4% tax on our cash flow as an added tax and they gave us a sheet of paper saying this is what you're not supposed to do. We were able to send anything through our system except for telephony or anything that was in direct competition with the PTT. But as any sensible person knows, there's no possible way to restrict that unless you fiddle with the ports on the routers and do some real technical stuff to stop it. But neither I nor the government can really stop people from using the Internet as a telephone. If someone has a credit card, they're going to use it. They use Vocaltec software to talk to their friends in the US and it works. So even though the government has written on paper that, No, you can't do it-it's the same thing as in India-people do it.

COOK Report: I get the impression that your cash flow situation is on a curve that will support increased bandwidth, but how do you match the cash flow to the customer demand and bandwidth? How do you keep all three of these elements in balance?

Agrawal: It's not easy. We have some real problems with the PTT. For example, say

I get a 128Kb/s circuit. They claim they don't have the equipment to get 128kb/s from the earth station to the local exchange. They have a 2Mb/s connection and 64Kb/s time slots. So they say we can give you either 64Kb/s or 2Mb/s. They don't have the equipment to multiplex anything else in between. This is where INET Access mail list comes in. The solution is very simple. If we buy a multiplexor for them for both ends, it will cost \$4000 for each multiplexor. But it will solve the problem. But even then it involves going through a big bureaucracy which takes the issue all the way to the level of the Board. This makes it very difficult to get things done.

There are two issues here. One is the technical issue involving getting the equipment to the PTT. The second issue involves getting the pricing for it. They don't have pricing for 128. So how do they figure it out? If we ask for 128, they'll take the price of 64 and multiply it by two. They know how the pricing works because they deal with Intelsat. But they just take the 64 and multiply it by two. If a customer pays times two, that's good enough for them. So for us to get a 128 kilobit per second line means we'd have to get two separate 64 channels which means doubling our expenses. So it's something we have to think about. Can we really support another \$12,000 a month? That's what it comes down to. We really need 512 but we can't get it.

We are now in the process of getting another 64. We have to. We can support it because our customer-base is growing. And since we want to be able to provide some decent service, we have to increase our bandwidth.

COOK Report: What would be necessary to convince the government that they should change their ideas on the pricing of high level bandwidth? Is there any hope of doing this in the near future?

Agrawal: The PTT is planning to provide Internet services in the market sometime next year. There are a lot of politics involved in this. For now, the PTT has agreed to restrict themselves to providing only high speed connections and not dial-up services. So we would eventually get our connections through them. Probably by then they'd figure out how the pricing would work. But as I said before, they KNOW how the pricing works but they SAY we can only provide you with two 64s. If they had a multiplexor, they could provide us with 128. The problem is that they don't want to change their attitude because they know that the customer who really needs 128 will pay for two 64s.

COOK Report: Does the social-economic structure work in Nepal such that the government imposes surcharges in the business areas of technology in order to

support agriculture, education, health, etc.?

Agrawal: Yes. That's what the PTT tells us. It costs 150 rupees per minute (\$2) to call the US from Nepal, including taxes. So we can lower the price of our international calls. But that would benefit only about 5% of the population. So by overcharging you guys for your half circuit, etc., we're subsidizing businesses in our rural areas which don't generate any revenue.

COOK Report: It's the same question of universal access that the US is now dealing with as telecom becomes increasingly deregulated. The issue, even in the US is becoming more controversial.

Agrawal: Yes. Who is going to take care of the rural people? It's more difficult and expensive to lay cable in rural areas, for example, and the returns aren't good enough. And to some extent, this line of reasoning is valid. But because we're in business, we don't want to hear that.

COOK Report: Where do you see things going in the next 2 or 3 years? Do you feel you have a reasonably stable growth curve that will enable you to continue doing what you're doing but on a larger scale? What changes do you think you might need to make?

Agrawal: Once the PTT comes out with their Internet service, we won't have to go through them to get 64Kb/s channels. Hopefully, we'll be able to get a 512 through them at a very reasonable rate. At that point in time, the PTT will have to do something about their rates. They say they will give a 2Mb/s line to some big name.

COOK Report: Sounds like it would be a win-win situation for you to send them to TeleGlobe.

Agrawal: We tried to send them to TeleGlobe but they're looking for a Sprint or MCI. You see, here they don't know who TeleGlobe or UUNET are. Sprint, AT&T and MCI have name recognition.

The North American Fiber Highway System

"This year, America's Network has requested the assistance of KMI Corp. a research company based in Newport, R.I., to help develop the magazine's Special Research Report on Fiber Deployment. Due to an exclusive agreement between America's Network and KMI, we are able to present readers with comprehensive map data of all North American fiber optic routes that are planned to be built or are already in place." From: <http://www.americasnetwork.com/issues/99issues/990215/990215map.htm>

ICANN Presses Forward: Broke but Buoyed by NTIA & GIP Intent to Salvage It

Editor's Note: As ICANN prepares for its Singapore meeting, we include some documents on its financial condition, on the efforts of NTIA to give IANA without justification to ICANN, and a flashback to October 1987 that is useful primarily for the insight it gives into the GIP as the likely master of ICANN.

Vint Cerf Passes the Tin Cup for a Broke ICANN

To: all ISOC Advisory Council Members
Subject: Funding ICANN From: Vint Cerf, Chairman of the Board, ISOC

It is now more important than ever before to provide funding for the newly-created Internet Corporation for Assigned Names and Numbers (ICANN). The organization is in operation. It has an interim board of directors chaired by Esther Dyson and an interim CEO, Mike Roberts. The organization has raised a modest amount of funding thus far, in the \$250,000 range, but it needs at least \$1,000,000 and perhaps as much as twice that, before it can establish a funding mechanism by way of the domain name registrars and internet address assignment organizations.

A number of companies have provided funds in modest amounts of \$25,000 to \$50,000 each. Contributions in any amount are, of course, welcomed, but plainly the financial support of all of us who benefit from the orderly management of these functions will be needed in this interim period while regular funding mechanisms are worked out.

The need is very great. Some organizations over which ICANN has cognizance are extremely well-funded. One domain name registrar has just gone into the capital markets to obtain a reported \$700M in a secondary stock offering. ICANN needs your support now just to be able to operate in the context of such well-funded constituents. As a non-profit, ICANN does not have access to capital funding of this kind.

I have copied Mike Nelson, who is managing a funding campaign for the Global Internet Project. I have also copied Mike (Roberts), the interim CEO of ICANN. Mike, can you do two things for us: 1. con-

Continued on page 21

Standards and the ITU Business Model

On September 5, 1998 **Tony Rutkowski** on the Cybertelecom list serve mentioned the major ITU meetings taking place in Minnesota this fall and complained that documents explaining what was at stake had to be purchased from the ITU

Dave Crocker replied: The ITU has existed for a long time with a particular funding model. The world is changing and the ITU is in the difficult position of having to make some changes, as indeed it already has. Groups within ITU can, and do, make specific documents available publicly such as when they are coordinated with the IETF.

Editor: this lead to a long and useful exchange about ITU policy

Rutkowski: As a matter of record, the ITU's assertion of copyright over CCITT standards and imposition of a significant charge and restricting dissemination, only occurred in the 1970s - rather late in an organization that traces its origins to 1850.

The money acquired was relatively minuscule compared to the overall costs of the standards process - and went not to the standards secretariat, but to fund favorite political projects. The effect has been to impair the utility of ITU-T materials - and covered in Carl Malamud's Internet Travelogue book. Carl has been a leader in encouraging the availability of public documents.

However, the original subject of my note on this list had to do with ITU Plenipotentiary documents potentially affecting the telecommunication and Internet policy and regulatory world. Those documents have not been made public. This has nothing to do with either funding or the IETF.

Crocker: Now, the market changes and one must adjust the business model. The discussion about changing business models is not in any way facilitated by calling the existing/previous business model a "boon-doggle". Quite the contrary.

Chris Savage: In my experience (I used to work for an ILEC), you either ignore such accusations as below your notice, or you take them on head-on and explain why they are based on a misunderstanding of the facts or the relevant policies. Which is ITU doing?

(Sept 8) Rutkowski: The problem is not with the business model, but the internal and external politics of the ITU environment and the policy choices that are made.

The preponderance of the ITU functions are

focused on "legislative" activities: producing radio and telecom standards (a/k/a recommendations); and a variety of treaty and other instruments. These produce enormous amounts of documents. Outgoing ITU-T Director Theo Irmer was fond of saying that the amount of paper would literally stretch from the earth to the moon. There's also a massive radio frequency registry function, and numerous different kinds of studies for developing countries and information collection functions for the telco systems of the world.

The numbers are frequently compounded by requirements that many materials automatically be translated into multiple languages because of general U.N. requirements. It makes no difference that there is zero demand for some of these documents in another language. They get translated at enormous collective cost.

All of this information resides on the ITU's information systems. The ITU actually does have what is probably the best MIS department of any U.N. agency. Everything is networked and in compatible formats. It's essentially zero marginal cost to make it available to the world. It's virtually all paid for by assessments levied against the world's governments. The ITU's current annual budget is about \$250 million, with an additional \$250-500 million or so contributed indirectly in time, resources, and domestic governmental activities worldwide to ITU forums. There are typically several ITU forums of various kinds occurring on any given day.

Only a very small percent of this material is publicly "published" - such as ITU-R and ITU-T standards. The ITU began copyrighting it and restricting its dissemination about 20 years ago to obtain an additional revenue stream for pet political projects - to cover the salary for positions that couldn't otherwise be approved through fiscal control mechanisms. The vast preponderance of information is simply not made available except to those who are formally participating in the activities. It's treated as "ITU internal" despite the fact that the ITU is a public intergovernmental organization.

Marty posed the key question - if these are public international legislative processes or other activities paid for by the public treasury, shouldn't everyone have access to them - particularly when the treaty instrument or standards are made mandatory in various ways? Shouldn't the U.S. and other progressive governments have required availability of ITU materials to allow public oversight? Wouldn't these ITU forums actually benefit through wider participation, contribution of ideas, greater awareness of the value of the work - particularly when there is zero marginal cost for their availability?

The materials remain restricted because all these activities are only nominally open, and the existing participants find significant value in keeping things that way. They don't want what are perceived as outsiders or even smaller players to access material. Even a player as large as MCI, paid its \$30,000 a year to participate in CCITT and hired a local firm to register for all the meetings, collect all the documents out of the pigeon hole, and ship them back to the States - just to get access to the documents and watch what the really big players and PTTs were doing. In addition, the notion of public oversight of the government is not a widely shared value within U.N. intergovernmental organizations.

Even as to the "published" materials like the standards, does it make sense to spend hundreds of millions of dollars a year out of public coffers and donated time of experts, then turn around and place impediments in the way of wide dissemination? Many attribute the triumph of TCP/IP over OSI in significant measure due to the open availability of TCP/IP standards. To use Carl Malamud's famous phrase - secret standards are not standards. This is probably where the ITU is most vulnerable - particularly in light of a dubious copyright assertion.

As the Internet effectively becomes the global telecommunications network, and the old environments merge, it will be important to assert "Internet values" of openness and public domain information. —tony

Savage: Even assuming an original (and maybe even ongoing) "political" motivation for making ITU documents hard and/or expensive to get, is there any sensible policy case to be made for doing something more or less along the lines that they are doing?

For instance, if the revenue streams are significant, one could argue "helpful" cross-subsidy — the money helps defray overhead costs that the member governments would otherwise have to pay. Conceivably one could argue that some of the documents are legitimately in the "don't try this at home, kids" category — stuff that every "real" industry player needs to know, but that, if truly freely available, would encourage hacking of the system, etc.

I'm just making this up, of course. My question is, has ITU asserted anything remotely plausible along these (or other) lines, or are they just sort of battenning down the hatches and telling everyone to go away?

William Drake This is the great irony of the ITU's approach, which Tony knows all too well and tried to encourage the rethinking of when he was there. The ITU may be shooting itself in the foot by restricting access to the standards set in ITU-T study

groups. Over the past decade, key standards-setting activities have been moving out of the ITU into other, more market-responsive fora, which has engendered a lot of soul searching and partial reforms in Geneva in the hopes of shoring up the ITU's self-declared "preeminence" in world-wide standardization. Asserting copyright and charging high fees for access may well limit the propensity of non-PTO players to design products based on ITU standards and hence help feed that decline, which ringing pronouncements at Plenary Assembly meetings aren't enough to change.

I think it would be in the ITU's long-term self-interest, as well as the interest of the entire industry, for the standards to be freely accessible and usable by all. But since ITU has such a hard time getting its members to pony up enough cash under the contributory unit system to fund all its activities (like using a high-priced Swiss printer to publish everything in multiple languages, resulting in huge stacks of unread paper sitting around in the basement), the ITU remains wedded to trying to squeeze bucks out documents sales, to its own detriment.

Years ago, before it had scaled up its own net capabilities, the ITU let Carl Malamud put the entire series of CCITT standards (@18,000 pages at the time) on the Internet for free ftp access on an experimental basis to see what happened. Thousands of people and organizations around the world who cannot participate in Geneva grabbed the documents and presumably consulted them when designing products. Having seen the possibilities of net distribution, the ITU promptly pulled the plug on Malamud's project, and rather than continuing with free distribution, moved toward fee-based net access. Wrong lesson learned. The tale's recounted in Malamud's book, *Exploring the Internet*.

(Sept 9) Raymund Werle: Carl [Malamud] and you Bill are right. Indeed, restricting the dissemination and use of standards does run counter to the idea of standardization.

On the other hand we know that those who have an interest in a standard usually do not hesitate to invest money and time to get a standard. Vendors send their well paid engineers to standards committees (i.e. they invest a lot of money) and most of them would also be willing to pay for a standards document.

Rutkowski: I think the point that Bill is making is that the more successful, market responsive forums and standards organizations typically found in the computer and Internet world almost invariably make the standards freely available. Where speed and mindshare are important - this is the only effective strategy.

There is a separate set of issues associated with public bodies funded at government expense and producing specifications that have a quasi-legal force and effect. There is a strong public policy argument that such specs should be made available at cost - which should be the nominal additional cost for making them publicly accessible on a server - which is essentially nothing - so small that the bureaucracy to establish an accounting system is far more than the cost.

ICANN Continued from page 19
firm or revise the information on where to wire funds or to send contributions 2. provide a list of companies from whom support has been obtained (assuming they have not requested anonymity) so that Internet Society Advisory Council members will know whether their companies have already made such a contribution. There is nothing more vital than to provide support for ICANN in this fluid period of formation. Your support is very much appreciated. Vint Cerf (2/12/99)

COOK Report: How sad that Vint Cerf as co creator of the most important protocol of the 20th century is now prepared to use ISOC as the means of installing an authoritarian unaccountable regulatory regime over that very Internet which he and Jon Postel played such an extraordinarily important role in creating.

I have interviewed Vint twice before and most recently opened the pages of the *COOK Report* to him less than a year ago by publishing his 4,000 word rebuttal of John Curran's analysis of the effect on the Internet of Worldcom's acquisition of MCI. He knows that I have treated him fairly in the past and will do so in the future. I have suggested that he explain in some detail why ICANN is so important and how he would justify accepting the kind of closed regime that ICANN wishes to impose on the internet. If ICANN can be justified by means of more than the brief hyperbole below, Vint ought to be able to do it. Regretably he has shown no interest in taking me up on my offer.

Quite frankly I find Vint's allusion to NSI in the note below not at all flattering. He should name NSI to make it clear that he is complaining about NSI not contributing to an organization (ICANN) that would like to dismantle it. It has long been fashionable to complain about the hated NSI monopoly. Now that we have seen how ICANN is ready to hobble the entire nascent DNS industry with a bureaucracy of which only the ITU could be proud, as well as prepared to give the trademark interests the upper hand over the business futures of small proprietors, I predict that we will soon be getting a "cure" far worse than the original disease.

Vint's mention of Mike Nelson's role is telling. I have observed Mike for nearly a decade during which he has for the entire time been devoted to promoting the interests of IBM. Those belonging to innovative green field companies who are being asked for contributions should ponder long and hard the meaning of IBM's promoting ICANN's interests. Given the current direction if ICANN is successful we are likely to find that a seat at the policy table will cost something along the lines of the \$65,000 charged by the WC3. Money and large corporate political power will be used by ICANN to further the interests of those who pay the most. The sweat equity entry fee of the IETF will be come a thing of the past along with American technology leadership in the Internet. When the bankers and lawyers move in those who live by their ability to innovate will find that Vint is using ISOC to destroy that ability.

That is the danger represented by ICANN and the reason why those who want to preserve the Internet as an engine of wealth creation and job creation should oppose Vint's plea and refuse to give ICANN a nickel for why would any rational person support something fundamentally hostile to his or her interests?

Major process issues are involved. I invite those who are interested to click on the following URL for a talk that I give in December at CANARIE to better understand why ICANN must be left to founder.

<http://www.canarie.ca/eng/networks/optical/pg15/cook-pg15.pdf>

That Vint would be blind to these process issues I find especially appalling, I also urge all ISOC trustees to debate long and hard before allowing Vint and his friend Don Heath to continue to use ISOC as the vehicle to, in effect, promote the global regulation of the Internet through ICANN. That it should now require two million dollars a year to do functions that basically have been done for a tiny fraction of that cost is absurd. Unless what he have going on here is an effort to provide ICANN with a legal defense fund war chest. Vint is inviting the community to deal with a two edged sword. The community needs to exercise great caution.

Gordon Cook, February 13, 1999

Fresh Insight from October 1997 News Story on Global Internet Project

On October 28, 1997 Rick Wesson posted a Reuters news story to the IETF list: Internet

Companies Welcome Idea of Global Charter BRUSSELS (Reuters) - A group of U.S., European and Japanese companies involved in the Internet informally welcomed a European Union proposal to draw up a charter to govern the global computer network. The companies, who have united as the Global Internet Project (GIP), said they wanted to be involved in the process, Peter F. Harter, global public policy counsel for Netscape Communications, told Reuters.

“(They) will individually provide input as to how industry may play the best role,” he said following a meeting in Brussels. EU Telecommunications Commissioner Martin Bangemann has proposed drawing up an international charter to deal with questions such as technical standards, illegal content, licenses, encryption and data privacy on the Internet and other electronic networks.

Watching the ICANN spectacle unfold in the context of the motives described by Peter Harter above causes me view GIP as the corporate designed instrument to carry out Bangemans wished for international charter. I wonder whether or not

1. the ICANN bylaws reflect precisely what those behind the GIP wanted. To think that Joe Sims and Jon Postel conspired alone and without outside input to cobble together such authoritarian bylaws defies credibility.
2. that major vetting of the trustees for the preliminary board was done by members unknown of the GIP
3. that changes to the by laws designed to keep the Europeans happy came from the GIPs EC contingent (witness the well known late November change to article 4 that opened the door wide to the ITU).
4. that Joe Sims was the hired gun to do the bidding of the GIPsters
5. that IBM's Roger Cochetti ran liaison between all the affected parties.

Gordon Cook Feb 17

Protest to the General Accounting Office of NIST Revised ICANN Sole Source Solicitation

David M. Walker February 24, 1999
Comptroller General,
General Accounting Office
Room 7100, 441 G Street NW
Washington, DC 20548

Dear Mr. Walker

As Editor and Publisher of the *COOK Report on Internet* I wish to protest NIST's Submission No. 296704 dated February 9 and posted to the NIST web pages on February 10. NIST's Amendment of SOLICITATION 52SBNB9C1020 is defective for the following reasons.

If NTIA wants to bridge the gap between the “anticipated expiration of the DARPA/USC contract and the anticipated completion date for the transition” it looks like what NTIA wants is the maintenance of the status quo - namely the continuation of the current authority where USC is the legal employer of the IANA people but where ICANN is paying the bill presented to it by USC each month. Someone needs to ask NIST to identify in straight forward terms what this solicitation is actually trying to accomplish.

Your auditors should ask how the “DARPA USC contract” could have only an anticipated rather than a firm date of completion? It must be a firm date and in the case of this contract (a temporary extension of the one that ended April 1, 1997) the date was October 1, 1998. Moving IANA informally to the DARPA/USC/TNT contract was a fig leaf designed to cover up the fact that ICANN has been since Oct 1 asking USC to remain the legal employer of the IANA staff, since it (ICANN) could not pay benefits and since USC could not legally pay their salaries from the DARPA TNT grant.

Since the contract is to be no cost, ICANN will have to continue to do what it is already doing - that is to pay the salaries regardless. What the solicitation actually does is put, retroactively, a legal face on the current status quo. The current status quo raises not only the issue of ISI's relationship with ICANN but also the issue of what statutory authority the NTIA is using to make these decisions in the first place.

NTIA is talking about the need for a contract to bridge from date A (the anticipated end of the DARPA USC contract) to date B (the anticipated beginning of full authority rule by ICANN). Unfortunately: (1) it can name NEITHER date; (2) the IANA salaries are being paid right now by the same entity that will pay them if the contract is awarded; (3) while NTIA alleges that achieving security and reliability in DNS management is critical to the stability of the internet, the fig leaf solution that it seeks by this sole source award (giving de jure validation to a defacto status quo) has only one value: - justification some months down the road of an NTIA decision to hand the root servers over to ICANN since ICANN then would already be successfully managing IANA on behalf of the US government.

Indeed the only argument having any po-

tential bearing on the stability of the Internet is the argument about the management of root servers for which NTIA has contracted already with Network solutions until October 1, 2000. If the stability of security and reliability of the Internet were in anyway dependent on ICANN, the major service providers would be taking pains to contribute to ICANN's funding. In point of fact they are doing nothing. Despite Vint Cerf's recent call for money on behalf of ICANN, roughly four months after ICANN's establishment, only just over \$200,000 has been contributed. This is roughly 20% of what NTIA said itself last November ICANN would need by June. Why should NTIA care more about the fate of ICANN that does the Internet industry?

NTIA's arguments about the critical role of the IANA staff are spurious because the only person who was doing serious IANA functions was Jon Postel who, unfortunately, is now deceased. The RFC editor (Joyce Reynolds) is being paid for by ISOC. Bill Manning was the only other IANA technical staffer with any visibility. But ICANN Interim President Mike Roberts told Manning back in September that ICANN would not pay his salary. The other staffers primarily provide administrative support for the computers and possess no unique knowledge. All significant decisions involving names, numbers and port assignments were made by Jon Postel. “dot US” (sometimes considered to be a part of IANA) is currently being paid for by the Postal Service there is nothing in the current situation that can justify NIST intervention. The GAO, in carrying out it's budgetary oversight responsibilities on behalf on Congress, needs to get specific identification of the IANA staff for which NTIA is seeking to transfer payment and supervisory authority to ICANN. In nearly every case those staffers either do not possess unique skills or are already being paid by ISOC. The dot US operation was merely a clerical operation inside the IANA offices that ran a registry. Absent Postel, it has no discretionary authority attached to it. It is a book keeping function. Nothing more.

Furthermore since DARPA announced the termination of the IANA contract on April 1, 1997 the question of whether there was ever any legal basis for the continuation of the IANA function at ISI or elsewhere has been in doubt. NTIA is trying to use NIST contracting authority to deliver a legal home to a package of functions currently being done and paid for by organizations as diverse as ISOC and the Postal Service. The only qualifications possessed by ICANN to become this legal home is the assertion of NTIA that ICANN does, in fact, possesses them. I am unaware that circular arguments have legal authority.

Executive Summary

Enron Network, pp. 1-10

We interview Stan Hanks who is Vice President, Research and Technology for Enron Communications. Enron is using power company rights of way to build its own fiber links from Portland Oregon to Salt Lake City, to Las Vegas, to Los Angeles and from Salt Lake City to Denver to Houston, Texas. Enron is using access to this fiber to do swaps with Frontier that will give Enron a national fiber network. Unlike Qwest and Level 3 and Williams, Enron is building a pure IP network. IP over DWDM without SONET and without ATM. This helps them to keep both capital and labor costs low by eliminating the purchase of ATM and SONET equipment as well as the need to employ engineers for the maintenance of SONET and ATM.

Enron's business model is very different from that of the other next-gen telco networks (Qwest, Level 3 etc.). Enron is offering an overlay of the public internet that can be used by large scale content aggregators such as Real Networks to deliver their content more cost effectively than possible by the purchase of multiple upstream connections and transit through congested NAPs.

In the case of Real Networks, Enron would have high-speed connections into their major distribution centers. The content that they originate and Enron distribute gets sent instantaneously to all servers at all of Enron's network pops. High-speed connections to ISPs, CLEC's and cable modem people run from these pop's.

The Enron business model is predicated on the assumption that it can function as a content aggregator getting content from producers content to local ISPs and hence to the audiences the content producers want to reach more efficiently and cost effectively through its network than through the public Internet.

In addition to this Enron offers a managed infrastructure where they handle all the network connections and routing and all the computing aspects of delivering the content. For example Enron drops a router and a server under its control and management into the pops of their local distribution partners. The router comes, at Enron's expense, with a high-speed leased line into the partner's pop. Enron then provides from its infrastructure to that of their partners all necessary high-speed interfaces. Enron sees itself as having a fix for many of the things that are broken about the current Internet traffic scenario. The fix is to move the more demanding content across Enron's transmission infrastructure. But note that Enron is not a general backbone network. ISP's can not buy access to the Internet via Enron. Enron is, in effect, a private overlay of the public Internet.

In the Real Networks model Enron is paid when the Real Networks servers are able to deliver requested flows via Enron's network rather than via the public internet. Last November Enron bought Modulus Technologies in order to be able to gain access to its Inter Agent real time control software. Enron content aggregator customers place this software on their web servers. When Inter Agent receives an http request it does a look up of the requestor. If that person can be reached on the Enron network, it ships the packets that way and writes the transaction to software that over time determines the total data delivered by Enron on behalf of its client.

Enron's overlay network is an interesting answer to the quality of service issue. In looking at the traditional IP class of service parameters you find that you have a very small number of bits to play with. Network operators have to deal with the dynamic tension between the requirements of operating a core network and the requirements of operating the intelligent network at the edge. They find that, at the edge, if they want to, they may offer very many flavors or gradations of Service. But, at the same time they will find that, as they take your bits across their backbone, the limits of technology require them to cut down those flavors and gradations to very small number.

While Enron has relatively few customers, it can certainly deal with QoS issues by applying a bandwidth solution to them. But when Enron to grows from dozens to hundreds, to ultimately perhaps thousands of customers, its ability to throw bandwidth at problems becomes much more elusive and it begins to need network engineering solutions.

IP Telephony Design Philosophy, pp. 11 - 16

In late January on the IETF IPTel Working Group mail list, a question by a 3Com employee about the SIP protocol kicked off a discussion that migrated from the technical aspects at hand to a very informative discussion of the general role of the IETF process in Voice over IP protocol development. The debate pitted those who want the implementation of an entire suite of standards that would yield a total voice over IP solution, versus those who are happy to see more rapidly evolved modular solutions that may be packaged to serve varying needs and environments even if they are all less than 100% interoperable.

The debate is also about accommodation versus innovation. Should we be striving to move the public switched telephone network to the Internet? Or should we use IP and computer technology to build an entirely new phone systems and set of services? The legacy phone companies and their mentor the ITU not surprisingly want a cautious conservative transition to globally interoperable standards at all levels. The engineers of the IETF favor innovation over predictability as they push for a process that can offer a mixture of computing and telephony services independent of POTs.

In a net head versus bell head display of experiment versus conservatism one of the discussants concluded: Give me choices in how to locate users. Give me choices in how to select IP/PSTN gateways. Let me choose which QoS choice I want (by giving ITSPs choices in products to deploy.) I want a *choice* in *what happens* when I pick up the phone to call grandma. I **do** want it to connect. I do want some base level functionality — just as HTML/SMTP/... all do their basic job on the 'net. But let extensibility roam free. If in the process products don't work with each other, fix 'em. Everything does **not** need to be perfect day-one. This has plenty of precedent in the IP world.

Nepal ISP, pp. 16 - 19

We interview Dileep Agrawal, an entrepreneur who has grown Kathmandu based World Link from his bedroom to a 50 person company in less than 3 years. Dileep shows the creative role portrayed by Teleglobe in offering reasonably priced satellite bandwidth downlinked to Copenhagen and sent by Teleglobe infrastructure across the Atlantic to US based connections. (One of Teleglobe's policies was a three month moratorium on his monthly \$4,000 half link bill in order to help him build up a cash flow.) He explains how the internet begins to play a role in the Nepali economy beyond that of mere tourism He shows how he must diplomatically work with the Nepali PTT in a country where a percentage of telecom revenue traditionally goes to support social programs something that makes cutting prices quite difficult.

ICANN Autocracy pp. 19, 21- 22, 24

ICANN continues its assault hiring a PR firm rather than opening its board meetings and getting ready to impose certification procedures for registrars such that Einar Stefferud commented on Feb 25: "ICANN threatens to destabilize all businesses that depend on stable DNS name arrangements." While ICANN appears ready to steamroller its agenda through at the Singapore meeting next week, it is also so broke that Vint Cerf publicly appealed to ISOC supporters on its behalf. Meanwhile WITSA, the international technology arm of the Information Technology Association of America sent out an astroturf brochure replete with factual inaccuracies asking its members to support the DNSO draft and the rights of trademark owners against the rights of ordinary domain name holders.. Finally we note an October 1997 Reuters article on the leaders of the EC and the GIP calling for the creation of the kind of international regulatory body for the Internet that ICANN seems intent on becoming.

ITU's Business Model, pp. 20 - 21

A discussion of the contrast between the ITU's and IETF's handling of standards documentation

Continued from page 22

When GAO demands that NTIA prove its assertions of ICANN's "unique" qualifications, it will find that ICANN is composed of a Board without expertise in the Internet, using bylaws written by an attorney without expertise in the Internet, and run from the bedroom of a retired "president" who has 25 years internet experience as an officer of an IBM founded organization of Internet using University chief information officers. Without Jon Postel, there is nothing whatsoever that is unique about ICANN or its qualifications. Therefore I am confident that you will find that the issuance by NIST of a sole source solicitation to ICANN is neither defensible nor legal. Sincerely, Gordon Cook

Finally on February 25, 1999 Einar Stefferud commenting on ICANN's posting of a series of detailed drafts and asking for comment with a 72 hour deadline wrote: I take the clear implication that this is designed to disable all useful discussion and comment. I also assume that all those who directly support ICANN are commenting privately to various or all BoD members, and that they will all be heavily counted against the public comments. Note that ICANN never has disclosed who those "many backers" are that formed what ICANN claimed to form an overwhelming consensus.

So, how can we possibly expect a different

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

story now. So, I predict that we will soon be informed that there is actually no consensus for OPENness, or for the PARIS Draft, or against the Registrar Accreditation plans (which are terrible from what I have seen)!

Just mark my words when all this is proudly announced by ICANN as being decided by overwhelming consensus support, without defining or identifying the parties counted in the consensus measurement.

I have never hesitated to criticize NSI for stuff that I disagree about, and I still think that their TM policy stinks! But, I see ICANN as a much greater peril to the whole future of the DNS, and I see that what ICANN is doing is more likely to perpetuate the NSI monopoly than anything else, as ICANN appears to be planning to impose the same NSI/NTIA invented business regulation model on all DNS registries and registrars. It is totally clear that ICANN has made no effort to disclaim this charge, and until ICANN does disclaim it, I will claim that they do in fact have exactly this in mind.

The problem is in the Market Structure Failure which ICANN policy promises to perpetuate by limiting new TLDs and imposing heavy regulation on the entire DNS industry. ICANN policy appears to be focused on forming a global DNS CARTEL, while NSI is welcoming competition from other TLDS in an open market.

**IP Insurgency: Internet
Infrastructure & Trans-
formation of Telecomm
- \$495 Now available.**

Subscription Rates

Choice of either ascii or Adobe Acrobat (PDF) format 1. Individual; College or University Department; or Library; or Small Corporation - \$225 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