The COOK Report on Internet

Broadband Spread Spectrum Wireless Extends Internet Reach of ISPs & Field Research Scientists

Increased Radio Speed and Decline in Price Enabling Smaller ISPs to Compete with Cable and DSL Internet Access

NSF PI Dave Hughes Explains How Wireless Data Gathering by Environmental Scientists Will Yield Huge Increase in Network Traffic

Volume IX, No. 4, July, 2000
ISSN 1071 - 6327

Editor's Note: Dave Hughes is the owner of Old Colorado City Communications. See <http://www.oldcolo.com and manages, and for his NSF research projects — http://wireless.oldcolo.com>. He has been an independent networking visionary for more than twenty years. He was technical advisor to Big Sky Telegraph, linking remote schools of Montana to a University and the net in the late 1980s and early 90s. He owns his own ISP company, and uses wireless in its both up and downstream connections. Since 1995 he has been doing research for the National Science Foundation in the grass roots use of no license wireless in remote areas, including Mongolia, and beginning in the fall of 1999 no license wireless and satellite technology for biological and environmental science. We interviewed Dave on March 18, 2000. We also updated the interview in early May, 2000.

Part One
Wireless ISPs

COOK Report: It is said that there have been tremendous improvements in spread spectrum radios in the last couple of years. Are these mainly driven by better digital signal processors or by a whole lot of stuff?

Hughes: It's a combination of things. One thing that kick-started the whole rush to go from two megabit radios, (the E1 standard in Europe), to 11 megabits per second was the Harris Semiconductor 'Prism' chipset which had a much faster chipping rate and permitted the development of much faster direct sequence spread spectrum radios.

COOK Report: What is a chipping rate as opposed to a CPU rate?

Hughes: It's how fast the chipset can handle switching from one frequency to another while handling the packets of data being broadcast or received. Harris explanation is at http://www.zettweb.com/CDROMs/cdrom006/prism/

COOK Report: I see, in other words, how rapidly it can dance around.

Hughes: In a manner of speaking, yes. With direct sequence modulation, under the FCC rules, with a faster chipping rate, the FCC was able to permit the basic signal processes that allowed manufacturers across the board to jump up to a rate of 11 megabits per second from the previous 2Mbps.

COOK Report: In other words, if you get a 500% increase in your chipping rate, then you can get a 500% increase in your bandwidth?

Hughes: It's not quite that simple.

COOK Report: But it's related?

Broadband Comes to Spread Spectrum Wireless

Hughes: Yeah, it's related to the rate, and of course you realize, when you say 11 megabits per second, that you're also saying LAN. Wireless LAN, 10BaseT protocol speed. So there's no accident that everybody came roaring out at 11mbps because you will find a 10BaseT connection built into the back of all these radios. And with such a connection, you end up extending the wired LAN from inside of a building, to a wireless wide area network, or WAN. That connection can span a city, or reach across a rural county, at the same speed as if the network were all wired. Acting just like a wired network. That's a huge, huge step, because for the first time these radios operating at LAN speed can connect separate LANs without slow, or costly, wired pipes acting as a bottleneck for their performance.

Now I wasn't too impressed with Aironet radios three years ago, when I was buying radios for the first NSF project. But now they have come out with an 11 megabit per second family of radios, as did Lucent, as did Solectek and others. They all happen to be running, though, under the FCC rules for 'direct sequence' modulation, which is a different way to modulate the signal than frequency hopping. Now Breezecom, a very successful company that started in Israel and has been brought to the U.S., also has frequency-hopping radios that, for the moment, operate at less than LAN speed. They are good radios and have a large market share.

COOK Report: Okay, now do both of these "flavors" (direct sequence and frequency-hopping) fall under FCC Part 15 no license regulations?

On the Inside:

Hughes: Oh, yes.

COOK Report: From an operational point of view why would you use direct sequence rather than frequency hoping or vice-versa?

Hughes: They differ in their ability to handle interference. Direct Sequence Spread Spectrum radios have more efficient switching by using Phase Shift Keying (PSK). They seem to be cheaper to make. And since they can operate at 11Mpbs while Frequency hopping radios are limited to 2Mbps, that is a huge advantage. The rules of the FCC to this date, have not permitted frequency hopping to reach 11 megabits per second.

COOK Report: And still do not?

Hughes: There's pressure and formal filings now at the FCC to change their rules to accommodate the new technologies. For example take Breezecom, which has a large piece of the market. They're quite successful, and we've used their radios, too. They operate at 1 megabit per second for 7 miles. The problem is they're frequency-hoppers. Given what their competition is doing, they want to come out with frequency hopping radios that operate at 11 megabits per second. And as soon as the FCC approves the changes - and there is no reason to believe they won't, then frequency hoppers can get up to the same speed, in the same frequency bands that everybody else does. But all that underscores just how bleeding edge these technologies are - the FCC is having a hard time keeping up with the rate of change.

COOK Report: What are the prices of the 11 megabit per second radios?

Hughes: They have started coming down. The early 11 Mbps Solectek Radio was $9,000 two years ago. Then out came Aironets, Lucents - competitor's that started out thousands of dollars cheaper. Right now, I can buy a PCMIA card for my laptop that cost $165 from Aironet that operates at 11 megabits per second, half duplex. Which really means about 5-6 Mbps true throughput. Which is 3-5 times faster than a T-1. And I can buy a Teletronics 2.4Ghz, 2Mbps radio for $500 for the base unit, and only $99 for the client radio.

COOK Report: So is it safe to assume that you can buy an actual 11 megabit per second radio for under $1,000?

Hughes: Yes.

COOK Report: Way under a thousand?

Hughes: Yes. But there is a difference between buying a Base Station radio - that can handle hundreds of client radios, and a one IP/MAC address radio - for the client end. Both prices have come down rather dramatically, with client radios now below the magic 'consumer' $500 price point. Which happens to be what I had to pay for a 1200 baud Hayes Micromodem back 15 years ago.

Now let me expand on what I have been saying. This is a technical speed-up. It is one that's been coupled by more companies jumping into the game. With more learned about the software/hardware required, and with more production. Significantly, companies are now getting gobbled up by bigger companies. This is happening as these larger companies — this includes the Ciscos and Nokia's and so on —have started to get into the wireless world. And some is unlicensed, but some is licensed, like MMDS and LMDS. Which opens the door to far cheaper 'backbone' IP links.

COOK Report: Craig McCaw has huge slices of that stuff.

Hughes: Yes, because you can make radios — again, with very fast signal processors and so on — that operate in licensed spectrum, but get high rates of reliable, error corrected, reasonably secure, data communications. Up to 100Mbps.

COOK Report: And they're up into the high number of gigahertz frequency, right?

Hughes: It depends. MCI Worldcom bought up a whole lot of MMDS licenses. And they've come out of the 2.5 gigahertz to 2.7 range. That's not way up high. That's just above 2.4, which is unlicensed. And they're heading towards services that are still aimed at the corporate or business level and that get up into quite high speeds. Fixed, wireless networks.

Cisco Integrating Wireless and Wireline

Now here comes another recent development. Very, very significant. Cisco, a device-based company, manufacturer of routers par excellence ends up doing several things. It buys Aironet for $880 million. This is all in the last 60 days. They're getting into the wireless LAN market at that level. The necessary radios are down below $1,500 and with PCMIA cards down below $200 in cost. Aironet has made some changes since then, however. They have dropped the power of their radios to 30 miliwatts which is more appropriate for wireless LANs than ISPs. While you could still stick an amplifier on the antenna, doing so raises your costs. They have momentarily ceased production presumably prior to getting revved up again in their Cisco corporate guise.

COOK Report: Do the Aironet radios speak TCP/IP perfectly well?

Hughes: Oh, yes, perfectly well. Most of these spread spectrum no license radios are transparent to the protocol moving over them. But also, at the same time, Cisco itself comes out and announces an LMDS radio that'll operate 30 miles, cost $20,000 each end, but goes DS3, 45 megabits per second 30 times faster than a T-1. So there's a radio in the licensed area and that radio is out in the marketplace right now. The point is that the DS3 LMDS radio is integrated into and part of an extension of their wireline routers. Lots of integration is done in software. They're software experts. They can use the crunch capability of their router knowledge and software knowledge to drive the hardware that's already released. I have recently been briefed by them. They have some very interesting things up their sleeve. And they are not a communications service company, they are a company which makes devices - that communicate. If Cisco does things right it could represent a big paradigm shift from the Telco Model of business.

Now, remember, there are the no-licensed spread spectrum bands around 915 megahertz, and 2.4 and 5.8 gigahertz. But the FCC also came out with the UNII band.

COOK Report: That was about two years ago?

Hughes: Closer to 3 years. Now there are radios being made also in the 5.8 gigahertz UNII no-license bands that do not require spread spectrum.

COOK Report: At a 5.8 gig?

Hughes: It's the same thing. They overlap, as a matter of fact. In other words, the two band's services overlap, but that's normal in the FCC Part 15 frequency bands.

COOK Report: Aren't they severely artificially limited in their range?

Hughes: You put your finger on it, I've been screaming about this from the day it started. And you could only use one full watt of power in a UNII radio if you built a radio that offered 20 megabits of throughput. But then the FCC, always fretting about potential interference, imposed a formula for spectral density, whereby if you choose to make a radio that's only 10 megabits per second, or 5 or 1.544 - T-1 speed in the UNII bands, you must lower the power accordingly. And so every way you do it, the range of a UNII radio is not going to exceed five miles. The laws of physics still hold true. And that's okay for a lot of uses. But its lousy for anything really rural, or even school districts which are spread off across cities. Yet the UNII band was touted as having 'solved' the connectivity problem for and between public schools. My NSF studies show that no license radios need to have ranges of be-

tween 10 and 15 miles if they are to serve even 90-95% of all public school districts. Not 5 miles.

Cisco is bringing out UNII radios. They're going to be pretty low cost at the client end - perhaps from $500 to $750 at the beginning, and they're going to have other people, like Motorola manufacture the end users' radios, while they manufacture the main base engine at home, in Cisco. So my problem, which few others have expressed, is that the UNII technical radio manufacturing restrictions the FCC imposed are too limiting for rural use, where interference is not the main problem, distance of link is. The very places the President is citing as having a huge 'digital divide' - affordable rural connectivity to the Internet.

COOK Report: Clarify for me what you meant when you said $500 to $750 on the client end. In other words, a base station is, what, point to multi-point, which is more expensive and more costly than the end user device?

Hughes: Yes the base station is much more sophisticated, of course. You could even take a PCMIA card and have it talk to another PCMIA card. Period. But the minute you put a second card, two cards talking to one, then you have to have what's called a base radio, or an access point. That's essentially the generic term for it. An access point is the point at which multiple radios talk to one radio and then presumably out the back of that access radio, your traffic goes into some network by either 10BaseT or 100Mbps Ethernet. Or, as the way LANs work, by moving traffic back out via radio to another computer on the LAN, through standard hubs.

In the Aironet line, for example, is the BR500, which 11 megabits-per-second radio and operates in the 2.4-2.483 gigahertz band. It can theoretically have over a thousand connections to it. Of course you are sharing its bandwidth. But that is no different from a hundred workstations on a wired LAN sharing a T-1 pipe to the Internet. Except this radio operates much faster, at least 3-5 times the throughput of a T-1. The BR500 radio itself lists at $2,400. But the end user radios, which can be a PCMIA card, can be as low as $165. And with a 'pigtail' that permits an outside antenna. One client radio talking to the point-to-multi-point radio. So there is usually a base radio - an access point - and a number of client radios, usually called bridges. They differ in price.

Now, another one of the advances that have been made is that that BR500 not only can be an access point radio, it can be configured as a relay radio. You don't have to have a different radio, as you did with the early BreezeComs. Meaning if you can't reach your destination because of non-line of sight problems or distance, you can put one BR500 at the base, one BR500 out ten miles and another one another ten miles or over the mountain. You can have 20, 30, 40 miles with one or two 'hops' using that same radio at every point. Moreover, that same radio can be logged into by password protected telnet or a web browser. They can have IP addresses in them. Even a crude router. You can use a web tool not to just configure, but also to monitor the wireless network, remotely, even if the radios are in a box high on an outside tower.

COOK Report: In other words, the newer radios come with a user-friendly GUI interface that will allow you to set it up in your own network of multiple radios?

Hughes: That's correct. I use these myself, in my own ISP operation. I originally replaced my T-1 connection to my upstream provider, 3 miles away in downtown Colorado Springs, with three Breezecom 3mbps half duplex (making them effectively the same as T-1) radios. Three radios because there is a large office building blocking a good line of sight between our building in Old Colorado City and a 14 story office building downtown Colorado Springs, where a number of ISP backbone providers are. So I set up a relay point, which, because Breezecom's Access Point/Bridge radio designs, I had to use the AP radio at the relay point. That was fine for 2 years, rain or shine. But then, as things grew and newer radios came out, I replaced the Breezecoms with 11 mbps Aironets, and simply 'configured' one of them as a relay radio.

Now, I can log into that radio that's up the street or up a tower. I don't have to climb the tower. I can log into it, I can monitor it. And when Qwest had an outage yesterday for four hours, I was able to ping all three of my radios and then finally the Cisco router, all of which responded that there was nothing wrong on my end, caused by a wireless outage. So when I called Qwest's network operation center, I knew the problem was theirs, not mine.

Then, inside my ISP offices, we have the usual wired LAN network between our servers and router. And we have a number of business customers in the same building where we are. These customers wanted higher speed connections to the net than dial up. They can't get DSL, don't like ISDN, and would have to pay, each of them, over $1,500 a month for a US West-MCI/Sprint/Quest network connections. We merely ran 10BaseT connections to them through ceiling tiles, and sell them a fast connection from $35 a month for just the link to the outside, to $55 a month for that plus full ISP service - email, web space, shell or net accounts on our server.

Now, since I want to be connected from my home-office, we have another set of radios, from Wi-Lan of Canada, 915Mhz frequency hoppers that deliver T-1 speed to me in my house. Free. And to the History Center, and its web site. And to another home-office. Then inside my house I use an older generation 'only' one mbps wireless LAN to my laptops. For use anywhere in or around my house, like out on my porch, smelling the flowers. While running my business. And connecting at the lowest speed of one mbps.

But here is another point. The Breezecoms are still in use! For a business with several workstations, in another building close by, wanted a minimum one Mbps connection. I installed the Breezecoms to link the two buildings to my net. So now I have four different radio brands connected by hubs and 10BaseT cabling to the same Wide Area Network in Old Colorado City. The moral of the story is that radios can be reused easily in different places. Lots different from yanking out wires or fiber.

And I am able, sitting at home, on a wireless LAN, to log into my Aironet radios, check the status, see how many packets it's sending, what the error rates are, or change the radio rates, set up separate wireless channels for groups of radios to operate in, within the total allocated spectrum, permitting sub or parallel nets. Check for retries, error correcting, messaging, and many other configurations. In short, all but the very lowest cost, and earlier - 3 years ago - spread spectrum radios have gotten very 'smart.' The era of the smart, self configuring radio, has arrived.

So the large companies are buying up smaller, successful wireless companies. And that's giving the whole industry credibility. And that's also what's getting the interest of Wall Street. Because when a Cisco or a Lucent and so on starts to buy up smaller com-

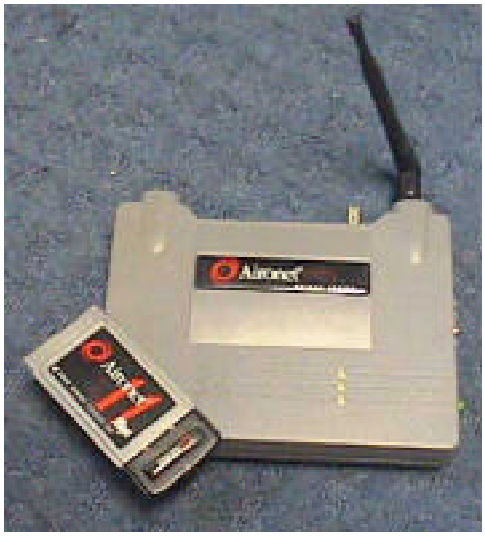

Aironet BR500 11mbps Bridge, and PC C ard Direct Sequence Spread Spectrum

panies that are successful, it's a very clear signal then they're getting into the wireless game. Whether it's licensed services on a monthly basis or whether it's unlicensed. Radically different from Telephone companies buying them up, to just put them into the tired old central office, per minute cost, circuit switched telephone business model.

This is all, of course, a relatively recent phenomenon. And while there's something in the neighborhood of 7,000 ISP's in the U.S., the fact is that hundreds of them now are delivering wireless services.

Wireless ISPs

COOK Report: Well, I heard you estimate about 300?

Hughes: That's just an estimate. There was a hard number of 180 months and months ago, but the increase is evident. It may be a thousand by now. Nobody's counting. I am watching it on a daily basis on the wireless ISP mail list. [Editor's Note: the list address is isp-wireless@isp-wireless.com and as we went to press with this article we were told that Breezecom, alone, now claims to be selling to 500 ISPs in North America].

I have been reading the wireless list for a long time, I'm seeing, on an almost daily basis, somebody new to the list saying, well, I'm an ISP and I want to offer wireless service. Or I want to start offering next month wireless ISP service from scratch. What do I need to know?

COOK Report: And they're talking a wireless service meaning connecting to an upstream wireless or allowing users to connect to them via wireless?

Hughes: No, both. Starting with users. And some of those have started out, of course, by aiming at business. And government and schools and so on.

COOK Report: It becomes another way to crack the local loop.

Hughes: Oh, definitely. It is absolutely a local loop issue. A large number of them are in pretty awkward locations. They're in rural areas where we're talking about towers for the radios.

COOK Report: And right now, we are talking local loop broadband, which as of two years ago, we weren't

Hughes: That is exactly right. Right now you have a lot of people looking at WaveSpan. UNII band. You can buy their Stratum 20 radio that goes 20 megabits per second UNII for about five miles and they maintain it can go more than that. [**Editor's Note**: Wave Span was acquired by Proxim in December 1999.] But, you can upgrade that radio to a Stratum 100 at 100 megabits per second. There are a couple of other companies which claim products in the area as well.

Well, that's significant enough that I immediately referred the San Diego Supercomputer center folks to that. Because they had an immediate, obvious need for it. The Cooperative Association for Internet Data Analysis (CAIDA) has to move its offices a mile away from the high speed network links at the Center. CAIDA needs high bandwidth for its network research. PacBell, could only talk to them maybe about a DS3, but you damn well know what telco DS3's cost. And two would be horrendous for small operation needs really fast bandwidth.

So the 100 meg radios are here. Most of them in the UNII band. Thus limited. But the FCC is also opening its mind, because there's pressure now being put on from many, many directions. In fact, there's proceedings taking place right now in which the widening of more bandwidth even at the 2.4 gigahertz, is a possibility. And there are big companies and it's controversial. Because a lot of people think the little guy's going to get squashed. But with Cisco in the game we are talking now about even bigger companies. And there are now proposals at the FCC for ultra-wide bands.

COOK Report: Ultra-wide band meaning?

Hughes: Ultra-wide band means wider than the current no-license bands that are permitted for various no-license radios. In fact, there is an ultra-wide band — you got to be real careful of the terms here — that means low power and ultra-wide to some. And no less than the engineering staff of Paul Allen's Interval Corporation, filed with the FCC. Their filing basically recommended that the band cut across all bands. And cut across television and FAA and everything else. [Editor's note: on April 21, 2000 Paul Allen announced through his Vulcan Ventures holding company that he has closed Interval Research Corp after eight years of pursuing advanced research. Approximately 30 of the staff will be offered an opportunity to join a newly formed Allen venture, which will focus purely on advanced development for Vulcan's broadband-oriented portfolio companies.]

COOK Report: Because spread spectrum is so good that it can just, it can do its thing and not interfere with others?

Hughes: They used Paul Shepherd's MIT thesis in which he mathematically demonstrated millions of radios in Manhattan, not interfering with each other, and passing hundred's of megabits per second, with radios, architecture, and of course FCC rules, to permit it. Signal processing, it isn't just spread spectrum, spread spectrum's almost a passe name now. But the idea is still, across many frequencies, sharing with other emitters and receivers, very low power - sometimes below the 'background noise' level in the area, and digitally processed.

So there are a whole series of proposals at the FCC taking place and some are being brought in by some pretty substantial companies. For example, the MMDF by MCI Worldcom, who spent $400 million getting frequencies in the 2.5 to 2.7 giga hertz range. Their Warped One would be a 1.5 megabit, a T-1, for $300 to $600 a month. But the 310 would be, called Warp 310, $40 a month for 300 KBS. MCI Worldcom, using licensed MMDS is talking about consumer level now and about direct competition with DSL and cable.

COOK Report: You have the continual slow movement of the RBOC's on DSL and, I think, the growing perception that DSL service is just going to degrade, the more DSL's you have on your copper loop from the central office. The Bell-headed telcos not doing a terrific job of rolling DSL out. The current AT&T/TCI strategy and now the AOL-Time Warner strategy for broadband access via cable TV is bogged down in controversy. Consequently, it begins to look like the shine on both DSL and cable TV as a broadband mechanism for TCP/IP is looking rather tarnished.

Hughes: You're absolutely right. Remember that DSL, is based upon the central office locations and an investment of about $150,000 to get a DSLAM into each central office. This is according to figures that I have heard directly from COVAD. And then, regardless of the overload, you start out running into another problem. As you get out toward Edge City, Suburbia, the size of the market starts dropping. And then you get into rural and small towns. Well, when it really starts dropping, then there's a crossover point where DSL is just not economic. Most of rural America will just never see it.

COOK Report: If you're looking at wireless for the lower cost issue or the issue of more or less immediate competition to cable and DSL for broadband into local loop, how important are line-of-sight issues with some of those radios?

Hughes: It's very important. However, it's interesting that Cisco, as part of their marketing, even on the frequency they're using for LMDS, maintains that they have developed technologies that limit severely the line of sight problem. In other words, diminish the problem technically.

Line of Sight and the Fresnel Zone

And one of them is called the Fresnel Zone. The Fresnel Zone is the fact that Radio waves when they travel from radio A to radio B, travel in kind of an elliptical envelope. So depending on the frequency, depending on the range, you have to have so much clearance above intervening obstacles, or from the ground at where your radio antenna is located.

I'll give a very, very specific example that I was involved with, down in the San Luis Valley. There was a tower that was down close to San Luis, which was down in a depression, with a ridge up above it. From the top of that ridge — and from San Luis to the top of the ridge is only a half a mile, maximum. But from the ridge top it's 30 miles line-of-sight to downtown Alamosa. It's a clean laser-light shot there.

The tower that we wanted to use was down on the shoulder of the ridge just above San Luis but only peeping over the top of the ridge to Alamosa, 30 miles away. By having a surveyor go up on the ridge and checking it, we found out that there was only ten feet of clearance from the top of the tower looking over the ridge line-of-sight to Alamosa. Well, it was line-of-sight from the top of the tower and it was line-of-sight down ten feet, but no further.

Fresnel Zone calculations for that radio at that range and at that frequency say you want 50 feet of clearance. And as a consequence, when, getting up on the tower, we put up a standard radio there, we found that it would not connect, because a big chunk of its power was diffused by hitting the ground before it went all the way. And yet, standing on the ridge at ground level with the same radio, with a handheld yagi antenna, with the same power, where the ground dropped away in front of you - obviously down about 50 feet, we got a good, 30 mile connection. Consequently the only way to get around that was to go to higher power. Because we were an NSF funded experiment, we were able to get a waiver to do higher power. According to Cisco, they are now able to handle up to 49% blockage by the Fresnel Zone and still get through. So, technically, they say they have basically overcome some of the line-of-sight problems.

COOK Report: But with some of these radios, if you're talking about using them to get a broadband connection from residential homes to an ISP, you presumably would have to have that ISP radio attached to an antenna on the roof of its building. You would also quite possibly have to have an antenna on the roof of your house? In other words, if there's trees in the way, that's going to hurt also?

Hughes: You cannot generalize about things like trees. Because it all is a function of frequency, effective power and range and the nature of the obstacles. And that cannot be generalized about. So one of the fundamental realities of wireless, no matter what it is and where it is, is you must have a professional site survey before you buy and install the radios. That is unavoidable.

COOK Report: That's still the case, even with the better interfaces, and so on?

Hughes: Yes, because you're basically measuring RF signal strength Fresnel Zone obstacles. The trees, deciduous trees are not the same thing as pine trees. Wet trees are not the same thing as dry trees. Dense trees are not the same thing as sparse trees. Trees with snow hanging on the branches are not the same thing as branches in a snowstorm - when a properly sited radio will work just fine. Both through snow and rain.

COOK Report: That web site picture with the great gobs of snow on your trees, by the way, was a good visual. That picture was worth 10,000 words.

Hughes: That, of course, is it. It even blocked my little half a mile signal that's six-tenths of a watt from a Canadian Wi-Lan Radio. A good radio that is up 99.9% of the time, rain or shine over the past year.

COOK Report: That snow on the trees was like pulling down a window shade between you and the ISP.

Hughes: Yes, that's why unless you use a professional installer, you better learn a lot about how radio waves work. It's not rocket science. Certainly no harder than configuring Routers and Servers. Just different. And if distances are short, and you have good radios, you can do lots of rule of thumb work. The Aironet radio has software that permits you to walk around with a laptop and see, visually, the margin of connectivity power you have between points, with different antennas. Its the long reaches, when people try to get that last mile, out of 25, that takes skill.

Special Skills Needed by the Wireless ISPs

Hughes: Exactly right. And so you have that. So there are companies... in fact, a great deal of the discussion on the ISP list is about what does it take in distance and so on with given radios and how do you get around it? There's a guy in Texas who is superb on the subject of installations. His name is Jaime Solorza. 915-778-5966. And ask for Jaime.

The point is, they are installers. They are resellers. They are site surveyors. They do an enormous amount of business and they really know this stuff from top to bottom. And you can ask him any of these questions, but also you can even ask him for some good examples that he would know of there.

COOK Report: For example, I'm getting the impression that if you've got an ISP in a city of a couple of hundred thousand population and it's sort of semi-rural you have to have specialists. Or even if your ISP is in a city of 500 or a million population, that ISP either has to have, among its own employees or with a partnership company, someone who can do this kind of survey if it wants to advertise for wireless customers.

Hughes: Yes, on the initial set up. But it depends upon the skills of the ISP.

COOK Report: In what sense?

Hughes: Now, who are ISP's? Well, these are people who can handle Linux or Sun machines and IP and routers and so on. You've got to have that in your organization. But the ISP technical people must also learn the radio work. Siting and installation doesn't take an RF Engineering degree, but it does take some rigorous stuff, so there's lots of talk about spectrum analyzers and so on and rules of the thumb. I'm experienced enough so that I can take a look at your set-up and I can usually tell you whether it'll work or not. And if not what it would take to make it work. Know who the best friends of local ISPs can be? Ham radio operators. They may know little spread spectrum well - even though there is an entire branch of Hams, TAPR who specialize in digital radio - but they sure know antennas, and wave propagation and reception. And even local laws and ordinances about placing antennas and towers or masts.

While the U.S. rule is a maximum of one watt of power at the radio (much of Europe limits it to 100 milli-watts) and then there are rules for maximum permitted antenna gain, what's called the EIRP, the effected radiated power. And then there's a whole science in antennas. And there are rules about how much total effective radiation you can have, radio power plus antenna gain. One watt at 902megahertz is going to go a lot further and through a lot more interference than at 2.4 gigahertz with the same power. A good example was the Mongolian installation that my CO-PI Dewayne Hendricks did. Because the Mongolian Engineers hadn't answered our questions about buildings in Ulaanbaatar, and what you could see from their downtown building roof - we asked for a video tape - before Dwayne arrived in Mongolia, we weren't sure whether we were going to have outside antennas on every one of those seven sites

Well, it turns out the Russians didn't put steel bars in a damn lot of those concrete buildings. And so up to about 3 kilometers, there was no requirement for an external antenna, because the 902mhz signal went through the walls to the little rubber duck antennas on the 1 watt. FreeWave serial port radios. Had there been metal in those walls, that wouldn't have been the case. And we had to guess when we shipped thousands of dollars worth of associated gear and antennas besides the radios. We would be there for just 10 days - no time for later shipment. Because we were experienced and had a little luck, we guessed right.

But this doesn't always mean that you've got to go up to the roof to get more range than a rubber duck antenna can give you. It can also mean that you have a little flat antenna just outside the window. Or a directional antenna, even inside the building, that concentrates the power and so may make the difference between success and failure. One of my NSF projects was 'Local History by Wireless' in which I put an NT web server in the computer room of an old church building, now a museum for an historical society, and connected it up 3/4 of a mile - through trees and even a big brick building - wirelessly <http://history.oldcolo.com>.

An inside ground floor antenna would not connect. An outside yagi antenna on the roof of an historically restored church would look incongruous. So I put in the attic, above the insulation blankets, but directional. An 18dbi gain yagi. It worked, and is still working 2 years later. Radios are as much an art as a science.

So there's lots of ways to adjust things, but that's part of the site survey. It's experimentation, or it is on a big network, it's really a professional thing. And the larger companies, like Solectek if you're going to buy a whole network for a company in a town that may have six branch offices around town will send out an RF engineer, who first of all does a path analysis with software. I have a piece of software that costs $1,100. It basically takes into account the terrain, from U.S.G.S map sheets, range and then applies frequency calculations. It will give you the calculations that get you really close. From that point on, it's Reality Time, and you do what needs to be done on site.

COOK Report: So of the 300 to 500 wireless ISP's, are most of them connecting mainly a few small businesses?

Hughes: Yes, most ISPs offering wireless, offer it to businesses first. So they can learn from it, among other things. And expand from there. Until Teletronics came out with their WLAN product line, with end user radios as low as $100, the radio cost was too stiff for residential connectivity. And as these companies learned, by putting little connectors on even the PC Card radios, so external antennas could be connected, it becomes more and more possible to connect economically to home owners, and profitable for the ISP. At rates equal to or above dedicated telco services, or DSL, or cable. Bypassing them all.

There's a whole range of companies including a large one in Utah that specializes in providing high speed connections to large businesses, in direct competition with U.S. West. In many places wireless is the *only* direct competition with the telcos. I wish Congress and the FCC understood that. They flap their lips about 'competition' and about 'the digital divide' - then virtually ignore a whole emerging industry under their noses, while trying to regulate older technologies.

Wireless Business Models

COOK Report: What are the wireless ISP business models? Three flavors perhaps? ISP's that are using wireless to get to their upstream, might be one; ISP's that are using wireless to connect businesses and business offices to each other and to the ISP in place of leased lines would be another. And then is there a third where ISP's, if you want to have a radio in your home and you're just an individual user, who are beginning to do that, that's probably the most recent and the fewest?

Hughes: Yes, but that's not a very good characterization. Because, there's a totally different problem going upstream than going downstream. If you're in a big city, there's very little advantage, if you're doing really high bandwidth to be going wireless up to that upstream ISP. The real opportunity is to be found wireless in delivery that last mile or that last five miles. Or between two small towns, or suburbs.

So you can't really generalize on it. But for small towns, it's a killer, because your cost of a small town ISP is not simply getting down to your customers, your cost is also getting up to the larger city and wireless becomes very significant there — it's called microwave. And microwave is up to 100 megabits per second now. And microwave costs have come down. And so licensed microwave, which work pretty damn well when you're pushing this 50 mile range, is sometimes the method of choice to get from your town to a larger city.

COOK Report: But if you're 20 or 10 miles or something like that and you can go line of sight, then there's some other high speed radios that are pretty good.

Hughes: Right. Cisco has an LMDS radio that can do up to 30 miles, line of sight at 45 megabits. Also Jaime Solorza would be a good one to ask this question, because remember the thing I got into with Texas versus the FCC. That the FCC, as you know, has this god awful rule, involving the e-rate. Because the rule in e-rate is that the school cannot, or the library, cannot own the equipment. That is they cannot buy either the microwave equipment or even a satellite ground station (they're never cheap) or a pair of radios using e-rate funds. And provide their own connection between their building, or to the upstream ISP. Dumbest damn decision the FCC ever made.

COOK Report: That's just absurd.

Hughes: Well, it's absurd because it's expensive. And it's absurd because it's forced the schools into an annual recurring cost contractual arrangement with telcos, even though it's subsidized cost. If and when the Congress decides to kill that program, they're all going to be standing there holding the bag with their infrastructure built around that annual cost set up. The telephone companies are laughing all the way to the bank.

Now, today, for example, in the San Luis Valley is a good case, the 30-40 or more miles in the rural areas still have a $2,000 a month local loop charge for even a T-1. And with yet the radios now exist that down in the San Luis Valley, we could go on top of that ridge and we could easily be going five megabits per second true throughput for that 30 mile distance and down to that school. Doing this would cut out, totally, the $2,000 a month local loop bill for that school district. And the district would then pay only for the five megabit per second Internet bandwidth cost from the ISP, or the ISP and the school district could agree to 'choke down' their bandwidth to say, T-1, and pay less. Like maybe $1,200 a month. Now, its $2,000 a month to the telco PLUS $1,200 to the upstream ISP. Or $3,200 a month to the school district that is 40 miles away from the big city. Buy a pair, or even three, $500 2mbps radios, which incur a one time cost of less than one month of telco connectivity (which still requires that you buy a DSU to connect to them), you suddenly are down to $1,200 a month. That's the comparative economics of wireless.

And remember also that many ISP's have awakened to this, partly as a result of being approached by their customer, where the customer owns a pair of radios. The customer provides the extension from the business to the premises of the ISP, not the other way around. It's a very smart thing to do, because then you own a pair, you could take them wherever you want. And all you're doing is getting permission from the ISP or the building owner that they're in. That, of

course, is what I did with Colorado Supernet 3 years ago — permission to put the antenna in there and the radio in there. Because of this 10 Base T general method, you're handing the ISP an RJ45 to connect directly into his router. He's not paying the phone company to come back to you! So the upstream Internet carrier is paying less too! Consequently, they charged me less.

So there is that absurd FCC level decision on the e-rate. You better believe I had that on my mind when I was invited to speak before the Texas Infrastructure Fund which is a Texas version of the national E-rate. And it was set up on the same principle. It comes out of rate-payers pockets and goes into a $1 billion, ten-year fund. Managed at the state level. And it was for subsidy of schools, libraries and, in the Texas case particularly, health centers.

All right. The rule was, before I made my speeches to them, that you had to have a service. If you were in Cut and Shoot, Texas, 40 miles outside of Houston, you had to use the phone company. And if that 40 miles cost $2,000 a month, that's it. You'd get the subsidy out of that fund of $2,000 a month until the end of time. And then once you're inside Cut and Shoot, there's the separate issue of how you distribute the T-1 signal, even between buildings of a company or buildings of a health center, like clinics, or the library or the school or all of the above. More monthly telephone line costs.

So I went down there to Austin, and showed the Infrastructure Fund what it could do with wireless. I said, for god's sakes, if you Texans have any brains, you'll change your rules from the way Washington and the FCC does it. They changed the rule. They may be Texans, but they aren't rubes when it comes to dollars and cents. And it's a very nice formula. The formula is, since it's based upon the phone company subsidy, essentially, if a T-1 from the ILEC or CLEC, is going to cost $2,000 a month, they normalize it to T-1, to Cut and Shoot, or from wherever. You multiply the monthly figure by 12 and you get $24,000. And you add to it any cost of equipment that you'd have to have under the T-1. For example a DSU, CSU, or whatever.

You cowboys in Cut and Shoot may now apply all of that to an alternate means of connectivity. And it can be no license wireless, it can be licensed MSDS wireless, it could be microwave. And $24,000 will buy one hell of a lot of radio delivered bandwidth, both to the town and within it. Even relay points and everything else. And then you're only left paying for the upstream ISP. But as you know, an upstream ISP in many cases is not necessarily a commercial ISP. It may be a government. It may be a university. That's a very common thing.

As a result in Texas, any school, library, or health center, now has an alternative to continue paying monthly-telco bills with the sole future prospect of rate hikes and unending dependence on the telcos.. Would that Washington would get that smart. But then one has the feeling that the politicians are so cozy with the Telcos, they can't see over their shoulder to the Wireless Future, which is gaining on them fast.

Flexibility from Extreme Low Power and Antennas

There's one other thing I'd better say in between this. In the 2.4 gigahertz area almost everybody makes radios that only operate with 100 milli-watts of power. One-tenth of one watt, or one-tenth the authorized power by the FCC.

COOK Report: And why do they make at only a tenth of a watt?

Hughes: Because it's cheaper. It's much costlier to make a one watt radio. And they can sell the same radio in Europe, where the rule is 100 milli-watts everywhere. So they make their base radio that way because there are now many companies making amplifiers to put on line between the back of the radio and the antenna. And YDI, Young Designs, in the D.C. area, is one of the best of them. And almost every company now makes amplifiers. You put the amplifier between the radio and the antenna. If a radio puts out 100 milli-watts, you add a 1/2 watt amplifier, you are still inside FCC rules, and you can get longer range. So, again, one way to solve the problem of connectivity range or interference by trees and so on, is if you can't get through at one-tenth of one watt, you then buy a plug-in amplifier to bring you up to a full watt.

Now, I have no amplifiers on my three systems. I don't need them. I get perfect connection, because the distance is not that great. I'm not fighting anything. But if I were fighting something, I would take the $750 Young Design's 2.4 gig one watt amplifier. You can buy up to a one watt amplifier on a 2.4 gig radio. And you're still inside of the FCC rules. So amplification is becoming a very popular solution to the difficult site, the long site, the one in which the radios, even with good antennas, don't do the best job.

COOK Report: But you still can't use an amplifier to exceed the FCC limits.

Hughes: That's right. But on the other hand, these radios are so damn good that you're talking about 20 miles with one-tenth of one watt. And a pair of the lowest cost radios have been measured at almost 70 miles with line amplifiers. All within FCC rule. It works. Its legal, It's reliable. And it's secure. And it's free, free, free. Who needs Ma Bell for bandwidth?

COOK Report: I hear you.

Hughes: Well, do you know how far some have gone? When I was in San Diego, I talked to Frank Vernon, a geophysicist who works with Scrips Oceanographic Institute. Vernon is a seismologist. He monitors earthquake data. He's got a huge bunch of three-way radios that are coming back up to towers and so on. But he flat said, out loud, in front of an audience of high end scientists, without amplification, he's got one watt FreeWave radios going 100 kilometers. I keep running across some pretty long stretches. 70 miles. 50 miles. Their performance is a function, within the rules, of the height, the clear air and, if necessary, and an amplifier in there.

Software Defined Smart Radios

COOK Report: So what are they doing at the FCC?

Hughes: There's a Notice of Inquiry out and it will be discussed next week at the recently formed FCC Telecommunciations Advisory Council - TAC. - The NOI is brand new and that is on the subject of a Software Defined Radio (SDR). One where smart software controls the radio - its power, its frequency spread, and other technical characteristics.

That's what Dewayne and I recommended back in April of '98 in our *Scientific American* article. We must permit the manufacture of smart radios which set their frequencies. And set their own power levels. Keeping the power to the minimum. It's a very, very important idea. We are moving away from the idea that radios have to be dumb and fixed and made for one set of emission rules. The radios now can be smart, intelligent and self-regulating, like the Internet. But that's another way to get not only more bang for the buck, but also to minimize interference in congested areas.

If you have smart radios, built to FCC specifications so the 'self-regulation' really works, then the FCC could raise the power rules! To 10, 20 watts! Then in the really rural areas, where distance is the problem, but where interference in those bands is minimal or non-existent, they could run full power - 20 watts, with higher gain antennas. 50, 100, 250 miles. But at shorter distances and in urban areas where interference is, or can be, a problem, the radios set themselves to, say, a quarter watt. Cause that's all they need! But it has to be approved by the regulators, the FCC, which is way behind the power curve on approving these new possibilities. Maybe the recent creation of the Telecommunications Advisory Council to the FCC

will help speed up change. My colleague Dewayne Hendricks is on it. And he is doing through his Dandian company on the island of Tonga, in the South Pacific. There, as they strive to link hundreds of islands, the Crown Prince sets the spectrum rules, not the FCC.

I submit the principle of 'smart radios' (and smart antennas) is a very fundamental answer to lots of 'scarce spectrum' issues nationally, and internationally, when coupled with digitally massaged data across wide bands of spectrum. George Glider talks about these things theoretically. We are doing them in the field, practically.

And remember that most of these radios also have sub channels that you can jump to. That's one way that they can it make so that everybody doesn't have to be in the same sub set. For example, in the Freewave radio, you can have 15 different settings within 902 to 928Mhz. And what does that mean? That means you can have this radio sitting here and communicating with a distant radio while you can place a radio right next to the first, operate it on the same general frequencies, and not have it interfere with the first radio.

Now, that can be set manually, of course, or they can even be set by being logged into. Now you're starting to talk about having about the little buggers scan their operating environment where they might find other Breezecoms in the area with some potential interference. Having done this, they make sure they do not interfere by setting themselves to operate in a different part of the authorized spectrum. And the FCC with the NOI (Notice of Inquiry, where it is asking 'the industry' to comment) has thrown the possibility of a world with such capabilities out there. There's going to be a debate, both technical and regulatory, because technology that operates under these premises makes it possible change the very way that the FCC regulates spectrum.

It's not just the dumb hardware of the past, grand fathered in forever, and it's not just the fact that it's no license and it's spread spectrum. But now we're getting into the area with the software defined radios, where, if the FCC is smart, they will also shorten the life of licenses. Manufacturers must upgrade their capabilities or lose their certification for their unsold radios. For we know greater capabilities are coming along in software, radios, modulation, and antenna design. We are in an era of accelerating progress of digital radio design and operation.

COOK Report: Well, who are doing some of the offerings of the smart radios?

Hughes: Ask Jaime. I just don't memorize all the makes and models. There are at least 80 companies now, of radio manufacturers. He's in touch with damn near every one of them and he would be able to answer that real quick. Proxim just bought up one of these outfits —Wavespan but Proxim was already in this game. Research is going both ways, it's going into more powerful radios, but it's also going down to miniature radios.

Now, I put out a question all over the place — what's the smallest radio in the world? By god, I got answers. I got a reference to a Dan Withers up near Seattle and the organization is called www.worldwireless.com. I am now able to buy, a Freewave for $1,250. Oh, easy, up to 115 kilobits per second, serial. One watt. Frequency hopping. Really good radio. Very, very useful. But, whoa. I just ordered from old Dan Withers a kit, i.e., two radios, which are 56 KB, one watt, almost everything else the same characteristics. Per radio the cost is $335. And the combination is $700 for the kit, including all kinds of stuff. And you might get up as much as $500, but the fact is the pair of radios now can be bought, a serial, that'll do 56 KB or 115. For lots of uses that is plenty fast - certainly for the environmental scientists I work with, whose $3,000 data loggers put out only 9,600 baud of data.

COOK Report: At what bandwidth or what range?

Hughes: Same range, it can go up to 30 miles or 40 miles. And being 902-928 megahertz, punch through walls. Now we're talking about an end user radio inside one's house, serial. At the 902 to 928 range, frequency hopping stuff. But then there is Teletronics and their low cost 2mbps radios.

COOK Report: Well, a minute ago you said you get what you pay for. Have you tested it yet?

Hughes: You get what you pay for in the company and the support. And the total corporate follow through, ease of configuration, good documentation, best possible performance. And all those little diagnostic features.

COOK Report: So when you get the cheaper thing, you may be a little bit more on your own.

Hughes: A little bit more on your own, a little bit fewer return phone calls. But a whole mail list to ask questions on.

Customer Driven Advancement

Hughes: Yes and this leads to the concept that the end users can connect among themselves and then one of a connected group can link to an ISP. It's not all downstream.

COOK Report: Okay, because what evidence is there that, if I know the capabilities of doing these various and sundry things, and I know there is a Sprint Earthlink POP is in Trenton, New Jersey, near my house for example, I could knock on the door and ask them to let me connect? With these big national systems, there's no way in hell, with the commitment of any reasonable amount of my time they would agree to connect my radio. But what you are saying is that if the owner of a small ISP has a POP that you can reach on a reasonable basis and is aware of what can be done, even if he doesn't have a radio program yet, you can call him up and say, can I come in? You see the question I'm asking.

Hughes: That's exactly what I did with Colorado Supernet. Giant MCI would not let me do it, even though they are in the same building. They didn't have, or understand digital radios. I'm the one who went to them. They blew me off. So I went with Supernet. Now MCI keeps calling with me, pleading with me to look at their upstream prices. I blow them off now.

COOK Report: But when you go to them, do you make the argument that you should cooperate with me because even at some retail price base hook-up, I am not occupying local loop infrastructure to get into and out of your pop?

Hughes: Sure. That's part of the argument. You know what the other argument is?

COOK Report: What?

Hughes: Consider the ISP as captive to the telephone company. If the upstream ISP cooperates with you, you're essentially showing him how he may go into the wireless business by using your equipment for starters. And learn what it does. It's a cheap way for an upstream ISP to get some experience and exposure.

COOK Report: But you said there's a physical device that you can bring to him that is the receiving radio that he plugs in where?

Hughes: Into a garden variety Ethernet hub! The back side of the radio has an Ethernet port. Let me talk you through this. In my house, I have this laptop. And it's got a PC card in it that's wireless. One megabit per second. It's older. It cost $650 four years ago. I can do it now at 2mbps for $100 today. The PC card radio talks to an access point. Well, what the hell is the access point? It's nothing but a little white box into which an identical radio is plugged, like the one that goes to the PC cards. (Did you know that the much touted Apple Airport wireless is nothing more than Lucent wireless LAN cards in Apple's box?) But the only thing the box does, it turns the radio signal into an

Ethernet signal. And it has a 10 Base T female connector there. That's plugged down to this little $59, five port hub. Ethernet is Ethernet.

So, between the two radios, it's radio protocol. But down to that hub, it is Ethernet. Now, coming out of the hub beneath my desk is another Ethernet that plugs into the back of this Wi-lan radio, from Canada. Which goes up to the roof to a Yagi antenna. Which is pointed towards my office. When it gets to my office wirelessly, it comes down into another radio made by the same company and set to be point to multi-point. Meaning it comes to me, but it also points down to my son David's house, so he's got a connection. And to the History Center's radio. Three of us share a T-1 connection. Could be 15 of us, in the neighborhood. Heck, one early adopter could set up an omni antenna on his house, serve 5, 10, 20 neighbors with a 2mbps or faster connection, then turn around and connect to the upstream ISP wirelessly, and split the cost 20 ways. It could be cheaper than DSL or Cable, and go where they can't or won't.

802.11 Interoperability

So I've got three different brands of radios and they're all connected and they operate at different speeds — 1 meg, T-1 and 10 megs per second. They're all normalized to an Ethernet. That's why it's not coincidental that the 10 meg is Ethernet speed. I could even go further. I could take one of these serial radios now and buy a little $40 connector that goes serial to Ethernet.

COOK Report: Go further?

Hughes: Meaning that it is not going to go Ethernet speed but it's going to go as fast as the serial port will let it. Either 56 KB or 115. But there is a demand for lower speed radios, especially for these scientists and a lot of things, there's still a lot of serial stuff around. What I'm trying to get across here is, that you have inter connectability. It's all an extension of the LAN. And, of course, the 100 megabits is not accidental. That's 100 megabit LAN. So when you come in, you come out the back of the radio into your premises, you have 100 megabit LAN, you better have the 100 megabits, across that room to your router and so on to go upstream.

COOK Report: As long as we're talking about this topology, I had heard that it is the 802.11 standard that enables the radios to interoperate. I had heard that if I have a good connection in my house, I could connect up my neighbors to my house and then to the ISP?

Hughes: Yes. You took the words right out of my mouth, because I was about to say, and this little $150, two megabit per second Teletronics radio will talk to the $2,000 BR500 Aironet radio. The 802.11 standard interconnects them! At 2mbps.

COOK Report: So if I had a $2,000 radio in my house, I could spread out a couple of dozen 2 megabit $100 radios throughout my neighborhood.

Hughes: You got exactly what the ISP's are doing.

COOK Report: Well, which ISP's?

Hughes: A whole bunch of them. For example Jason Simonds, Midcoast Wireless. 207-563-8080. See for example: http://www.midcoast.net/wirelessfaq.html ISP for Wireless ISP's he calls himself. Now, he is an ISP. And he is doing it. As I explained much earlier in this interview, I am doing it.

An ISP in Nome, Alaska, (www.nook.net) is doing it. He is operating a dial up ISP service in Nome, connected to the net via satellite. However, he took three FreeWave radios, attached one to his dial up server in Nome. He then placed a second radio with battery and solar power, as a relay way up on a high mountain ridge 45 miles northeast of Nome. Then 25 miles beyond the relay ridge lies the village of White Mountain. He took the third radio and attached it to a dial up serviver in the village. The relay radio on the ridge can see both Nome and Whitemountain. So he delivers a commercial 56kbps connection in Whitemountain, which gets to Nome for free, and goes from there to the net via satellite. He uses the radios to extend his connection at zero additionalcost to him and at a rate of 56 kbs from Nome to White Mountain where he has paying customers. And makes a profit!

Now, a lot of this stuff is still done by hackers. Like those who ran the early Internet. And the earliest computer bulletin boards. Before AOL. Remember them? There's a woman in the wireless ISP mail list who is just so ingenious, she just drives me up the wall. But the fact is, I've been watching her ask all these questions, she acts like she's a dumb blond. But by god, she's running the thing and it's working. And she's not super high tech, but once again, where is the expertise coming from?

From the mail list. E-mail. Talking to them. And there's expertise, obviously. You've got to do a lot of learning. Eight or nine years ago it was Ethernet. And routers. It was the whole evolution of the Internet as it migrated down. 5 years ago it was the Web, and HTML coding. Where the hell was the expertise in the early 1990s? Well, it was among the hackers at the bottom and they were on the ISP list. And they were talking routers. The point is, it's still in this — while it's serious stuff for a lot of ISP's and it's real business. The fact is there's a huge amount of innovation and entrepreneurship that's taking place in communications among wannabe wireless ISP's or ISP's that want to add wireless to their operations. And they are thumbing their noses at the telcos. And setting up services where no telco dares to go. Too unprofitable for such a behemoth.

COOK Report: Well, it's this same grass roots kind of stuff which served as the training ground for all the network engineers for the commercial services.

Hughes: You're exactly right. There is really two layers to this. Cisco and the rest of the corporate world is coming down into this. For fixed wireless stuff. Either servers or like Cisco, to sell the thing and do it as a turnkey kind of set up operation. And you're having these grassroots ISP's coming up from the bottom. And they're not very small. I mean, you talk to Jaime and just ask him straight ahead what's his annual billings. I got a hunch this guy's got a hell of a growing company.

Coming from the Bottom Up

And these are in-between guys, these are not the ISP's themselves. They are resellers. But they're oh so much more than resellers. They're distributors and they're resellers and they are themselves expert in this stuff. And they go out and do the site survey. And Jaime's answering questions right and left at the same time he asks some, because this stuff is exploding in many directions.

For example, there's the issue already of throttling. This means that the ISP can adjust your radio to give you only 256 KB which you pay for. Even though your radio is capable of 10 megabits per second.

This allows the ISP to price to his customers' needs. Everybody doesn't need ten megabits per second. They don't want to pay for it. It's just like any ISP. You're paying for bandwidth, right?

COOK Report: So are you saying if I'm a distributor of radios that I can make some modifications to them?

Hughes: Well, not modifications to the radio. You can use software that runs in the Linux system, for example. Or in the router. Add-ons that give you the capability to both track but also to monitor, but also to set a maximum flow rate to any given customer.

COOK Report: Okay, in other words, if I'm paying for an upstream bandwidth connection, if I have a radio that can come in and go 10 megabits and I want to put that on the guy's Ethernet, he doesn't want to give me

the capability to suck up 10 megabits from his system.

Hughes: That's right. Because if the radio is capable of delivering ten megs to you, and he doesn't want to let you have the ten megs, because if he does, you will put a service on it. You will resell your connection to your neighbors. I'm doing that in Old Colorado City Communications. I could, if I wanted, spread my wireless ISP business all over 500,000 population Colorado Springs. But I have to spend half my time showing biological and environmental scientists, from those around the San Diego Supercomputer Center, to one Hispanic researcher studying frogs on the top of Mount Toro in Puerto Rico.

COOK Report: So on a small scale, you've put in your own infrastructure.

Hughes: Oh, absolutely. And have had it up for 30 months. And part of it is true wireless from me. I'm using the wireless. But things like throttling down, the ability to do that in software, there's another thing. There is already, for Linux (there must be for bigger ones), it's called EPPP. Ethernet PPP. Now you know how PPP works?

COOK Report: Yes.

Hughes: The ISP has a block from a Class B, he's got a block of numbers. And he's got X number of telephone lines, say he's got 25 lines. So he has a block of maybe 30 IP addresses. But he has 8 clients or 8 customers for every one of his 25 lines. Well, why do you have DHCP and PPP? So that when you dial in, it temporarily assigns you an IP address.

But, now there's EPPP which permits you to do that over the Ethernet. So I could have a wireless based DHCP, because right now, and I've learned about that on this list, because I basically for $5 a month, renting IP addresses from my fund of 256. And I've got them to rent, but at some point, I will run out. And so basically by having EPPP with the wireless, I can preserve my supply.

COOK Report: In other words you can take a subset of them and you can multiplex them amongst a larger group of customers.

Hughes: Yes. And with bandwidth throttling of wireless, you basically can price and measure what you do all the way up and down the line. That accounts for a whole lot of the innovation that's going on.

COOK Report: So at the grassroots, everybody, you can develop a whole mesh of interconnectedness of everybody connecting to everybody else.

Hughes: You betcha. Now, here's the telephone number and the guy you want to talk to, Michael Young, YDI, Young Designs. He's in Falls Church, Virginia. Telephone 703-237-9108. And he sells radios. He's unhappy that I bought Teletronics radios and not all from him. His radios are good, but they're a higher price. And I've been there and done that. But I have bought his amplifiers. They are better than Teletronics. Now he's measured some of these other things.

But, for me what is important is the fact that I was actually able to buy these things and get them up and get them going between two systems just lickety-split. It's almost getting to be plug-and-play. And that's from a little PC card that cost me $99 and I only paid $400 for the access point. Which could talk to multiple radios, at 2 meg per second across the room.

And then go into the Ethernet hub. You see, that's the key. The key is that by going into either a 10 Base T Ethernet series or going into the 100 megabit level, you're plugged into a purely normal networking environment. There's no magical interface. And the radios can be modulated within that framework. So that's really what's been happening with wireless ISPs.

Part 2: The NSF Field Science Research

Tachyon, Globalstar and Qualcomm

COOK Report: OK. Tell me how the NSF Field Science Project ties into all of this?

Hughes: There are two studies underway. One that is very, very significant I haven't talked much about. That's the satellite delivery of IP.

COOK Report: That ties with Tachyon into your recent San Diego Supercomputer Center meeting, doesn't it?

Hughes: Well, it does tie it in, but Tachyon is just one of them. The generalization is that, IP delivered right down to the individual by satellite is coming on eventually, with lots of services. Well, that's, okay. However, an intermediate step is being able to come down to the point in a metro area, and I don't care if it's a neighborhood center or a to business or to a school system, at such a rate of bandwidth that it is practical to distribute the signal laterally by no license wireless. Now that is one hell of a model if you think about it.

In other words, everybody thinks that the upstream ISP has got to be someplace downtown. But what happens when you are able to hook up at 2 meg up, at 2 meg down into your little ISP operation, from a satellite, or your small business operation with four offices. Or with your school system. Or with your government office. And have your IP go straight to the net from a 1 meter dish aimed at a satellite, delivering standard IP packets. But then you reach your other offices or your clients laterally by wireless that also travels at a rate of at least 2 megabits per second. Out to 1, 5, 10, 20 miles.

Tachyon so far is the only one doing this and Tachyon has its critics. But if they deliver what they promise who cares? They are just beginning to attach customers. They state that the customer ground station that talks to the satellite is only $5,000 and you can do bi-directional, true TCP/IP, at 2 megabits down and 256k up, for $2000 a month, or 300kbps down and 64k up, for $795. This will include full IP services from any spot in Europe or the Western Hemisphere. They will plug your earth station into a terrestrial wireless ISP one of which is Concentric. There is the Tachyon business model.

But you have my business model, when you extend from that base station, out 20 miles in every direction, wirelessly, and split the cost of the monthly service between 20 clients. Because then this investment is economic in every small town in America. Suddenly you don't have any phone company involved at all. And that model is really significant, because that basically becomes a real solution for the most remote towns in the U.S.

Tachyon is one of the first satellite providers which does this inside a tolerable cost envelope. It is using the SatMex5 satellite system, launched in July 1999. Therefore it should be a while before the satellite wears out. See http://www.tachyon.net for more information. The Tachyon model of course also fits the most difficult, remote, field research. One ground station, on a hill, and 10 to 100 data loggers out in every direction - all linked to it, wirelessly.

COOK Report: Is the business model emerging that Concentric will offer a family of services via Tachyon for small, remote communities who can link into the Tachyon system and then from Tachyon to Concentric to the rest of the Internet?

Hughes: Yes, that's what they are trying to do. But it also fits field scientific research, which is why I am pursuing it. Globalstar is also a good bet. Now it went down in stock price when Iridium collapsed. But what's the huge difference? Iridium was analog. Globalstar, uses Qualcomm CDMA radios. And Quaalcom is coming out with their 2.4 mbs 'HDR' wireless technology. http://www.qualcomm.com/cda/tech/hdr/

COOK Report: Well, McCaw took a look at Iridium and passed on it, but McCaw I think is investing in Globalstar, isn't he?

Hughes: Don't know. But Globalstar, has the Qualcomm CDMA spread spectrum radio. It is basically only able to deliver right now 9600 baud. But, when I made a recent cold call to Qualcomm, they were so solicitous of me, that on the same day they delivered to me the cable I needed to both charge the damn thing and use the data at the same time. Consequently out in the sticks untended it could be getting power and transmit the scientific data that we want to gather. Before they had two separate cables. You had to manually shift them to go between data use and recharging.

COOK Report: This happened at the San Diego Supercomputer Center Wireless meeting?

Hughes: Yes. With the Qualcomm radio, as they issue it right now, you get a plug, there's a plug in the bottom of the radio and you go into the recharger. Charge the radio. You unplug that thing and then you put another plug in to do a data cable, RS232 to a computer. Two different plugs. But not two different sockets. There's only a single socket for the two plugs.

COOK Report: Yeah, you have to do one or the other, but you wanted to do both.

Hughes: Exactly. And for their techs, they were doing both. But for their business model, they didn't think of that. So in their lab they had the cables.

I actually made a cold call in the morning and talked to the business section. Didn't come on real strong, all I said is that I'm doing research for the National Science Foundation. I spent $2,600 on your stuff, including the car kit and everything else. $1,500 phone. And I said, But what I need is the cable that I understood before I bought it exists, but back at Qualcomm, not at Globalstar. I need it to hook up biological scientist's data loggers way out there where there is no cellular, no place close enough to link up with 20 mile terrestrial wireless, and of course, no telco or cableco.

And I got home at night at 9 o'clock and they had delivered to my hotel the cable that basically plugs in the bottom and has it halfway down the cable is a little plastic box that has an input to it. But much more significantly, I had a call back request to call the guy at home, the international marketing guy and when I did, he said, we'll brief you, we'll show you the next generation. I'll get that briefing in late July.

Satellite to No License Wireless Distribution

Qualcom is wireless. It is digital, not analog, as Iridium foolishly was. And it is a variation of spread spectrum. It's not free. But it's a lot less expensive than any other terrestrial solutions for really remote sites.

COOK Report: What kind of a satellite system is Tachyon using?

Hughes: They can use anybody's. They are not stuck to one. They did that intentionally. Now they are on SatMexV. They can spread their service by renting space on other birds.

COOK Report: So, in other words, they're really kind of an uplink, downlink infrastructure.

Hughes: Exactly. But bi-directional IP. That's very important. Other satellite operators are selling downlinks by satellite where the return to the Internet goes by phone lines. Of course if you are out in the wilderness with no phone this model doesn't do much good.

I talk a lot about this, because I think that's an integral part of the wireless revolution. It's wireless terrestrially, horizontally, and it's wireless vertically - to satellites. It's the combination that really makes it. I might just do this for kicks in my company. I'll get that satellite sitting on my roof here and I'll offer a separate ISP service to my neighbors, wireless. That model will work. I will get 64 KBS up and 300 KBS down for a total of about $795 a month, flat rate. $596 is Tachyon only. Tachyon plus the Concentric Internet connection is $795. Customers will normally purchase the service from Concentric — including the Tachyon ground station installation and four static IP addresses. Total throughput is measured. For example the total through put for the lowest priced service is three gigabytes per month. If customer exceeds this, he will pay 20 cents a megabyte for the extra data. The high end service allows ten gigabytes per month

COOK Report: So that pays for both the Tachyon prices and the Concentric link to the Internet.

Hughes: Yes. But what is somewhat significant on that one, it's like a telephone company demarc the Tachyon rep says, (where the phone company terminates at your premises). A demarc is where you plug into our ground station. What you do with it on the other side, it doesn't matter. The Tachyon cost is not one of these things, where if you use five computers, it costs you one thing, if it's ten, it's another. You're paying for bandwidth.

And that system — only because of the FCC — can go uplink and downlink at two megabits per second. The point is that Tachyon can offer a service at two megs now. Not what I'm getting, which is the low end 64 up and 300 KBS down. That's correct. But you see that neat little combination, because that opens the door. That opens the door not just to the U.S. That opens the door to the rest of the world.

National Environmental Observatory Network

COOK Report: Take us through a summary of the things that you saw at the San Diego meeting.

Hughes: I'll mention NEON, which everybody seems to know about. National Environmental Observatory Network. And the word 'observatory' is kind of key here, because what they're doing is gathering huge amounts of data from remote monitors.

COOK Report: So you're seeing a huge movement under foot in environmental science and in other parts of the science world to use wireless monitors.

Hughes: I would describe it as a sudden awareness that wireless is a big piece of the answer - data collection - to what they want to do.

COOK Report: And it's now economically feasible to do it and, if they get out and do it, it's going to be another huge input of bandwidth into the Internet?

Hughes: Yes. Because you must also understand the observatory concept here, the observatory means it isn't just the scientists getting data, it means that you and I can look at the damn thing. Everybody. Citizen science, said Larry Smarr. Meaning you can't afford to have high paid university researchers going out and getting all the data. This is a direct quote from him — you need to train 9-year-olds to collect data.

COOK Report: Because they want so much of it, it's so widespread?

Hughes: Because you have to. When you're talking environmental and ecological, you're dealing with a huge number of data points and all over and you've got a data collection problem that up till now has been a manual operation. Or a problem with limited resources. For example the federal government down in Puerto Rico, (the Forest Service) is out there with more expensive stuff than even the colleges use. They showed me the Sutron data collector on a stream. Sutron especially sells to governments. It's not just better equipment, it's more pricey. But it's also designed to sit out

there forever on water courses, it isn't just for science engineering. It's also for monitoring flow rates, Army Corps of Engineer kinds of things. But then, it was going to a satellite, but in terms of cost effectiveness, they shut it off, it just wasn't justifiable, because it used to cost a lot to get data to the satellites.

What was NEON? The steps here is that NEON is a proposed project and they're having workshops, (this was the second) for which they bring in scientists. The focus is across disciplines from biological through environmental networking and computational. And the concept is interconnecting the scientists, the data and the databases and the visualizations and the standardization of data across disciplines. All of which requires a step up in data collection and reporting infrastructure.

What's the purpose of the meeting? The first part was for those of us who knew of technical capabilities, or in other words who knew what could be done, to be sure that the scientists sitting around there representing these various disciplines learned about the technology available to them. When they decide how the money ($100 million) should be spent, we also want to be certain that they don't think in terms of how they would have done it last year. Whether it is wireless, data bases, networking, visualization, or computation and number crunching.

Bringing the Scientists up to Wireless Speed

COOK Report: You're educating them about data gathering.

Hughes: Yes, how to gather data remotely, real time, and through the Internet right to the sensors themselves. Instead of gathering the data through data loggers manually, by making visits to where the data loggers sit in the wilderness. What became very clear was that I needed to rub elbows with more of these biological scientists than I'm seeing. With just two projects right now (Puerto Rico and Wisconsin) and I needed to see what they're doing in other places and other ways.

COOK Report: You saw it there big time.

Hughes: Well, in two ways I saw it. What I was totally unprepared for, because I was just going to be an observer in this thing, was how oblivious these people were to what's available right now or has even become available in the last couple of years. They are still in the 9600 baud, RS232, coming out of a little piece of equipment, manually connected by short cables world.

And they don't need a lot more, because a lot of the data gathering is nothing but a handful of numbers. But, so I showed them, across the board, all the stuff that you can do at higher speed. And I also went to Scripps.

COOK Report: Where you saw all the earthquake sensors.

Hughes: All the earthquake stuff. They were using a very elaborate set-up, all based upon the Freewave 115 KBS radio. But they could do it, because the radio can handle the data rates that they needed. And they were using Glenayre radios also for going to a couple of points. But I knew more about radio than they did. And they were still messing around with compression as a solution to getting more bandwidth through. That is an important point, because that's getting easier to do.

Then, I took him to Tachyon. And I also found that there was Mr. Wireless for the university. There was Mr. Infrastructure. There was from up in a different campus, Mr. Data Processing guy. These were the key central guys who run the systems. And there was a discussion about what they needed to do to get the data, but they kept mentioning how expensive satellite was and they just kind of ruled it out of their minds. And I said, right here under your nose is your answer. And then, they didn't even know.

COOK Report: And Tachyon's headquarters is San Diego, yes?

Hughes: Of course. They didn't even know that Tachyon has a ground station on top of the Supercomputer Center and it goes into their network at the San Diego NAP. And, of course, if it goes into their NAP, you don't have to go out by Concentric, right? If it's a research and education application.

But Frank had not seen that data. Then there was meeting of the scientists, where I made the second presentation. To hear that you can be doing 10 megabits per second with $500 radios just blew their mental doors off. Or when Frank, sitting in the back of the room, said, "at Scripps I go 100 kilometers with a pair of three-way radios." And I said, "Boosted?"

No! Right out of the box, he replied. FCC standard regulations. Anybody can do it.

Methods of Data Collection

Okay, so they saw that, but this whole data collection stuff is based upon sensors and entering devices and data loggers and data loggers from Sutron or Campbell — Campbell is one of the big ones — these things are boxes that are smart as hell. They cost from $2, 000-10,000. They sit out there hooked to devices, like weather stations, underwater sensors, light sensors, motion, wildlife sensors. And you can have many, many devices, hundreds of them. And then they collect it.

But in, almost every case, they collect it manually into a module which then can be detached from the data collector. The module, which just a memory storage device is brought back into the lab, which is at a forward research station and dumped into a computer with the software.

COOK Report: Well, that's the old way of doing it, right?

Hughes: Yes. That's the way that's very customary. Unless it's inside a lab. But this is the field stuff. And so that's the way they're doing it.

Hughes: Now, even Campbell does sell a connection to a traditional satellite transponder service. Big cost. Expensive way of doing things. They don't think about that. It's just too damn expensive. And anyway, your Internet's got to get to the forward research station, too. Which it doesn't do down in Puerto Rico. And the research station Internet link was only 56 KB when I went up to Wisconsin.

And here's another very concrete example — and I haven't got the solution, yet, but it's very typical. Right there in Madison, Wisconsin, the main university campus sits on Lake Medora. And the University has a Center of Limnology, which is the study of great water bodies. Now they have satellites pass over and taking very costly, scientific measurements, where the colors in a photograph represent temperatures and certain chemical properties of the water and so on. The problem is calibrating actual conditions on the lake with what the satellite sees.

So the point is that they've been sending people on boats out to some 60 different points on the lake. With a graduate student and on each boat, as the satellite passes over, the student grabs the temperature and a test tube full of the water. And they come back to the lab and analyze it. They then get from those points, temperature and water composition that they use to calibrate the colors on the satellite. They have no way to 'calibrate' the satellite, with real time lake data. To do it real time, instantaneously, on a mass basis. The methodology is a labor intensive use of many people in many boats.

Obviously, wireless comes in there, because, if one power boat came roaring around the lake and dropped off a little tiny buoy that had a radio which basically broadcast that data instantaneously, on command

or whatever, then that suddenly changes the nature of what they are doing. Paul Hanson, who is their chief tech, heard my pitch and ran out so fast they couldn't see straight. They now have $20,000 buoys and they bought Aironet radios. In order to see what kind of range the radio had, they put the thing in a boat to go all the way across the lake. They ran out of water before they ran out of radio. And so all of a sudden, that's big stuff to them. Then up in northern Wisconsin, their problems are not a matter of the labor on the data points, but in getting the data from instruments, situated out in the middle of a lake, on a raft.

Now, let's talk about the levels. Ned Fetcher is a researcher in Scranton, Pennsylvania, at the university. One of the things he's been doing down there is modeling the light on the bottom of the forest floor. There's two parts to what he does. He has a data logger. And then they have these little $15, sophisticated, photosensitive light diodes. His colleague puts out twelve at a time. They have about 40 of them out there at once.

Okay, but here's the problem. They only go out to about 25 meters, 75 feet. But they have to lay a wire on the floor of the forest. But people and animals and falling branches break the wire and they lose the data.

So he said, is there any way you can get a low-cost, $50 or less, radio. And so that's when I went to this smallest radio in the world outfit and I'm having a kit shipped to me, I think it will do the job. It has a chip on it and everything else. These cheap radios become themselves data collectors. Data just goes there and into a bigger radio, back to the center and then into the data logger and be processed.

In other words twelve or more of these little radios, each with a light sensor and a pattern on the floor of the forest will communicate back to a central point no more than 300 feet away. The data is sent into a radio, where it will be passed back to he field sation.

COOK Report: So the $50 radio goes about 300 feet.

Hughes: Right. And then the radio that doesn't even have to be an Ethernet radio, it can be a serial radio, like the $300 radio. That data gets collected, then, in the data logger back at the research station which may be a mile or two away.

New Technobiology Enables New Methodologies

So that is a way to do what he wants and he was really excited about that possibility. He

The Coqui Frog

said it would totally transform the way that he now has to work in order to get his modeling data. The experimenter down there, a woman scientist in Puerto Rico, Jill Thompson, said they have so much bad data because of broken wires that the experiment may be useless. And all of a sudden, if they can put that little thing out there without wires, then all you have worry about is theradio itself being stepped on.

But you've got another advantage here. It's real time. What really turned on everybody on was the ability to look at the data in real time. Not just to get it, because it goes back not just to the research station, it also goes into the Internet at the same time. And it may have to be wireless from that point back to the university. Because the research stations are usually out in the woods. And that's exactly what we're going to be doing in Puerto Rico. With wireless we will get from the field research station back to the main university and from that out to the 'Net, so that the researcher, who may be at any university in the world can not only see this stuff, but also see if something's gone wrong. And then theresearcher can talk to a graduate student who can go out there and remove the leaf that fell over the sensor or repair whatever has to be repaired.

Now I have covered two ways of researchers getting back wireless data. One is going directly from the sensor by wireless. The other is going from a data logger by wireless, which in turn is collecting stuff from sensors, that may be very close to it. The gathered data may then go either terrestrially with a relay back to some research center, or it may go directly to a satellite. And that's where the Qualcomm phone comes in. If 9600 baud's all you need, you simply plug it in and send your data back via satellite. Unless you need to send the data constantly, you can do it periodically, at a cost of $1.50 a minute. Maybe 4 minutes a day.

But let's look at the third level where you're getting into a little bit more. The Coqui frogs. Here we want sound. Not just some sound, but quality sound. A subspecies of Coqui frog that only live on the top of mount Toro. Requiring that researchers had to climb the mountain, after dark, in the nearly perpetual rain, go into a blind, record after midnight when the frogs sing, and then come down in the morning. Very labor intensive.

Why not use radio they asked me? Sure I said. Now I have to deliver. So we've got to have enough bandwidth to make sure this is not distorted. Enough that they can't, when they record it, back at the university, lose that all the flavor. Because I noticed that when we went there, and they asked me, can we go up this mountain to do this? They had even gone back to Sony to have them optimize the microphone to improve its reception in certain frequencies, to match the frog. And so the question becomes do you have to have enough bandwidth for fidelity of sound? Therefore a low end 56kbps serial radio is probably not enough.

But then the last project I was asked to do is pretty interesting because it is full motion video from down in Puerto Rico. The guy who wants me to do it heard me when I was down there last summer talking to 40 scien-

The Puerto Rican rain forest habitat of the coqui frog

tists. For eight years his work has been tracking freshwater shrimp under water. And he has underwater camera doing part of the data collection. He wants me to interface the underwater camera to a high enough bandwidth to have full motion video, which can be seen, back in the upper 48 states, real time. Not just have a graduate student, as he does now, go out and manually hold a waterproof camera down there, tape record the shrimp, and send the tapes by snail mail to the distant researcher.

Now, you're talking about higher speed radios, 5 to 10 meg radios or above. And yet that's cheap to do now. Distance is not great in this case and you're not going to go to satellite with that 10 megs. But you sure as hell can go back to your research center, you may have to compress, and do all sorts of other things, but nevertheless, the radios will permit this. They will go through the forward research station to the university, wirelessly, 15 miles away. From thence into the global Internet. I've already got cameras from Axis, a Swedish company (http://www.axis.com). The cameras themselves are web servers. With individual IP addresses inside the camera. With a serial port, and an Ethernet port. The Swedes only thought you could communicate from them by either slow cell phone, or in-building Ethernet. They never thought of 2mbps wireless radios connecting them up. However, I have. We are doing it. Not full motion, only 10 jpeg frames a second. But a step above still pictures, while below full motion with sound.

Watch out UUNET, the frogs and the shrimp are coming - using your bandwidth.

And so when I was holding up these radios in front of these biological scientists, they really jumped on me after it was over, and started asking me about coming to their projects. Including for example, one up in Michigan who uses parabolic dishes to capture sounds of both insects and bird life. And they would love to have those things sitting out there all the time. And not trying to keep a tape recorder going and all that stuff.

Tiny Linux, Forest Fires and Sensors

COOK Report: All right, can we finish up with the forest fires and the sensors?

Hughes: The tiny Linux thing, including Web, is the project that I have basically asked Steve Roberts, the hacker who used a recumbent bicycle in Mountain View, to do. Remember him?

COOK Report: The Internet-connected bicycle?

Hughes: Yeah, through radio. Then, but he now, he's working right now on canoes. He's got project going down the Missouri River and the Mississippi and going up the inland waterway and then back across Canada. It'll take two years.

Well, one of the things that he was playing with there and I jumped on when I heard about it, is the concept of a very small, solar-powered, mobile data sensor capture, database and communicator. With this you grab the data on the move, not with fixed points. So he's doing that right now on a subcontract, and we call it WANDER 2000. (Wireless Acquisition of Networked Data for Environmental Research). It's a prototype, which will be done by this summer. The device will be under 15 pounds total weight. It will have a miniature Linux running in RAM and a database in it and a variety of sensors can be hooked to it.

And so whether you're on the water moving or whether you're on a trail moving, with a backpack, horse, motorized vehicle, or whatever, you can take sensors and put it into the database, which can also be a website, using Apache, if you can reach it. And communicate it by either Globestar or whatever, and by other means.

But there's a very important point here. And it was reinforced at that meeting. You've got to have reliable capture of data even if all other communication is severed, such as during a hurricane. And so the small Linux comes into play, not just that it puts up a little website and a database, but it captures reliably the data even if your communications are down for one reason or another. It gets sent when you have re-established communications. The ability to cache in a data base. That is already mastered in Data Loggers. And in institutional computers. Now we must do it when linked, wirelessly. The tool to do this is a tiny Linux server, which may be in RAM.

And I was even asked is there any way I can go out on the plains of mid America where they have firestorms, prairie fires. Can I come up with a way to communicate the temperature of the fire at ground level and the gases that are being emitted, real time, while the fire is just feet away from the radio?

Challenge? Not for wireless, that's easy. But survival has to be dealt with. Okay, the combination of these things is what drove me to this wireless outfit and I'm really excited, because first of all, that prototype called WANDER 2000 will be done by this summer and we will basically exercise it. And the full plans of that, including the wireless connectivity to it, will be basically on the website, www.worldwireless.com. But at the same time, on that other island, Whidbey Island, is where this other guy is that I tracked down yesterday.

COOK Report: Whidby Island, near Seattle?

Hughes: Whidby, I guess that's where it is. But what happened is Dan Withers with World Wireless, whom I contacted, is not only a reseller for World Wireless' very tiny, low-cost stuff, super-miniature radios that include website capability. Tiny short range ones as well as one that's only going to be $300 to go a distance. He will be presenting his engineering at a national convention of sensor manufacturers. Because at that level at which he is operating we can actually open the door to transmission from the individual sensor. A weather sensor or something and not just a big complex testing device of some kind. Pushing the radio and the wireless data collection really down to a point.

And the way you do that is you don't expect to go all the way with that little radio, but you go to a next point of aggregation and the next point of aggregation. But it's not only going up to the 'Net, it is bi-directional. And so you can have a very tiny, specialized website with the data shown on it and remote access it. Note also that this fits in with the Globe project, the observatory idea.

COOK Report: The data is collected automatically and would go into these little Databases contained in the RAM of the Linux operating system on board the radio. And the sensors, then, are feeding into a radio within a few hundred yards or something? And as part and parcel of all this you automatically fed into a remote distributed series of Linux databases that do things with it there and then feed the data back upstream?

Hughes: Yes and you used a very significant statement there.

COOK Report: Distributed?

Hughes: Distributed. Because one of the things that was even discussed at the Neon meeting was how are you going to crunch

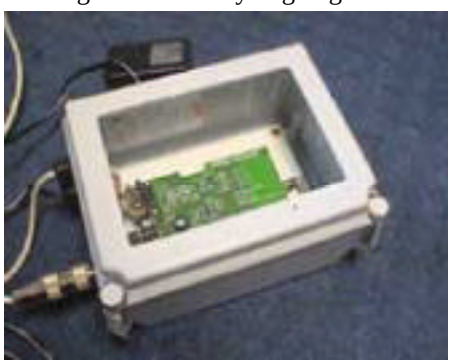

World Wireless Communications 900 SS 56kbps low cost radio, in field protection case.

all this data? One way is centralized, terraflop computers with high bandwidth between them. The other one is the distributed process. And, see, already, there's been a model for this. I can't give you the details on it, but there was a miniature Linux that went up that was in RAM with extremely low-powered electrical draw on a space shuttle that was used for the data processing and data collection for experiments that were on board.

The Tiny Linux has three capabilities. Well, first of all, it's IP to begin with. And number two, it can have a true database. And number three, it can process. Programs can process the stuff. And then number four, it can actually also be a web point, an accessible point. I mean, two-way, not just a broadcast way or not just a capture way. And the wireless connection from it gives you the bi-directional access to it. The only thing is, you've got to control that, you can't have 10,000 people all trying to look at the damn thing at once. So you have a management problem there. But the whole idea of the observatory is that you could call up a sensor sitting in the middle of a hurricane out on the tip of Manhattan and what's it saying right now?

COOK Report: A better example would be the estuary of the Mississippi River out into the Gulf of Mexico.

Hughes: But that's going to be, I think it's going to be almost equally true in the big, urban areas. Particulates and all that. Wireless, low-cost, no license, from low to pretty high bandwidth. Sophisticated sensors. And the ability to connect to and adjust or otherwise interoperate with the censor in real time.

You have the ability to use miniature Linux as a true IP handling device, and as an IP router. It can be a router from different sensors. Remember the conviction of those who say that everything in the world is going to have an IP number in it. In the summing up after they huddled up into groups and came back with reports, on the last afternoon and I said every damn Cocqui frog's going to have an IP number. We may be putting something around his neck or embedding in its ear. Or if we're really sophisticated, we'll read it out of his DNA. DNA as IP numbers. What a gasser!

But that's, of course, fundamental to the Internet. And it's fundamental to the IP flow that's there and the wireless just permits this in places that are inconceivable. And it's going to be extremely important to biologists and environmental scientists, because their problem is dealing with data and sensing and knowing what's going on in the most remote ways, way beyond where any commercial wire line is ever going to go.

Well, in closing, everything I am doing with wireless, from remote cabins in the mountains to frogs in the rainforest, is laying down the techniques for using wireless to every human being on this planet, wherever they are. And at data rates up to full motion real time video, affordably. There is a revolution coming for the Internet, thanks to terrestrial, no license wireless, digital signal processors, smart software, IP servers, satellites and the universal connectability of the entire global net.

ICANN Footnotes: What Some Others Are Saying about Arbitrary and Capricious Acts of ICANN, and Network Solutions

ICANN and Network Solutions now irrevocably wedded together continue to stumble forward. In a new series of ICANN footnotes we present some recent evaluations of their achievements.

Footnote 1

Beware of Monopolies Proposing to "Open Up" Markets: An Analysis of Network Solution's proposal for new top-level domain names

by Milton Mueller, Associate Professor, Syracuse University School of Information Studies (April 25, 2000) http://dcc.syr.edu/report.htm.

New top-level domains (TLDs) are badly needed, as the dot com space is getting increasingly crowded. But for five years changes in the TLD space have been stymied by political controversy.

On April 14 an official ICANN working group proposed to create six to ten (6-10) new top-level domains. The official working group report can be found at: http://www.dnso.org/wgroups/wg-c/Arc01/msg01095.html

On April 19, Network Solutions Inc. (NSI) released a proposal to ICANN to reduce the number of new top-level domains to two (2). Only one of the two proposed new domains (.shop) would provide an alternative to NSI's longstanding monopoly on registration in the .com, .net, and .org top-level domains. The other would be a restricted TLD for banks (.banc). NSI "generously" offered to operate the registry for .banc.

The NSI proposal is a step back from where ICANN should be going. It would slow the introduction of new TLDs down to a crawl and limit new domain name registries' ability to compete effectively with NSI. The proposal is designed to prolong NSI's dominance of the domain name market.

The NSI proposal can be characterized as profoundly anti-competitive for four reasons.

1. It would require the new (shop) registry to offer exactly the same terms and prices as the NSI com/net/org registry 2. It drastically limits the number of competing registries, for no good reason. 3. Its ownership arrangements would institutionalize cartel-like controls on the name space. 4. It would put NSI in charge of the back-office services of one the .banc registry, further reinforcing NSI's dominance of the domain name registry market.

1. The proposal eliminates competitive differentiation

The proposal would have ICANN sign a contract with a new registry "substantially identical" to NSI's current registry agreement with ICANN and the US Department of Commerce. That means that the new commercial registry would be forced to offer exactly the same terms and conditions, including price, that NSI now does. If new registries are unable to charge lower prices or to differentiate their terms of service, how can they engage in real competition with the well-known NSI dot com registry?

2. The proposal drastically limits the amount of competition.

The official ICANN working group charged to develop recommendations on new TLDs reached a broad consensus that there should be at least 6-10 new TLDs this year. This recommendation commanded a two-thirds consensus of the working group members, and was supported by public comments. The 6-10 number was proposed in order to achieve a more competitive marketplace and to allow a variety of different ideas and business models to be tested. However, NSI proposed to create only two new top-level domains. Only one of them (.shop) would be an open name space similar to .com/net/org. Thus, the level of competition created by the proposal is about as minimal as it can get.

The highly publicized Network Solutions proposal was part of a deliberate effort by NSI to divert attention from the Working Group's recommendations. At the Names Council meeting April 19, Network Solutions representative Roger Cochetti led a

Continued on page 22

Assessing the Current State of IP Telephony

Data and Voice Converging at the Protocol and Application Levels

Telephony Becomes Tool to Be Activated from a Web Page While New Web Oriented Applications Make QoS Less an Issue

Editors' Note: Jonathan **Rosenberg** is Chief Scientist for dynamicsoft Inc. and is responsible for guiding the strategic technology direction of dynamicsoft's suite of eConvergence products, which support the delivery of applications and services over next-generation networks. Before joining dynamicsoft, he was a member of the technical staff with Bell Labs Research, Lucent Technologies. He joined dynamicsoft in October 1999 and received his Bachelors and Masters degrees simultaneously from MIT in 1995. He is currently finishing his PhD at Columbia University.

COOK Report: What's been happening, I gather, since some point in the fall of 1999 to really change and speed the convergence picture between voice and data networks? One of the last things I published last summer was a very short essay that said, in effect, everything's kind of on hold at the moment in Voice over IP (VOIP) protocol development.

Now before that, during the Level 3 protocol development work that turned into Megaco, I published a really long, detailed, lengthy description of the protocol issues involved in getting the public switched networks to talk to voice over IP networks and so on. In view of that I would like to focus in more on the situation over the past year and assume a reasonable degree of knowledge on the part of my readers. Please help me to get a handle on recent changes.

Rosenberg: MGCP came onto the scene at about the same time that SIP was beginning to gain some sort of support. SIP actually has been complete for some time, I should say, with its own RFC in February of 1999. But I think people were trying to figure out what the story was and decide whether they were headed in the right direction.

COOK Report: If you go back a year, Level 3, with IPDC and then with Megaco, thought they had the universe by the tail and they thought that they had the key piece of the puzzle. However when you examine the strategy of the pre-standards group that they pulled together and the way that they hoped things would go from there, I gather that it didn't really go very smoothly in that direction.

Rosenberg: IETF and ITU efforts did merge although not quickly. But to the credit of everyone involved this is really a success in the sense that ITU and IETF both managed to work on this thing and both come out with the same thing.

COOK Report: And the same thing in this case is an offspring of Megaco?

Rosenberg: There were IPDC, SGCP, MGCP inputs to all these various standards bodies and both ITU and IETF decided to take it up. IETF started the Megaco working group, whose protocol was to be called, I guess, Megaco.

Megaco As Device Control Protocol

COOK Report: And that was December '98, I think, roughly.

Rosenberg: Yes, it sounds about right. And then ITU picked it up as well and decided that its output would be called H.248.

COOK Report: Okay, that's new information to me.

Rosenberg: So the agreement was to work those groups in parallel.

COOK Report: In 1999?

Rosenberg: I don't remember exactly when the decision was made, or which model they went for, but early '99, yes. They bounced stuff from IETF into ITU and from there back into IETF. This back-and-forth process was undertaken with the hope being that both sides could come to a document that was perhaps different in formatting, but identical in content.

COOK Report: So how did that work out?

Rosenberg: It did actually happen, but not without a great amount of air travel, stress, and bickering back and forth as you would expect.

COOK Report: So when did the baked goods come out of the oven and what were they called?

Rosenberg: Well, they're just cooling right now, actually. Megaco is not an RFC yet, but I think it has been submitted to the IESG for consideration as an RFC. I believe H. 248 is also going through its approval processes

COOK Report: So where will Megaco or H.248 fit into the big picture, now that they are finished?

Rosenberg: The protocols are identical, therefore let me refer to both as Megaco. Now Megaco functions as a device control protocol. The purpose is to take a large, telephony gateway, normally an SS7 gateway, and decompose it into three elements. First a signaling gateway which interfaces to the actual SS7 signaling messages. Then there is a media gateway that handles the actual circuits or the audio, both on the telephone network and on the IP side. Finally there is the controller, that talks to them both.

The deal was that, the media gateway has a fairly high amount of data volume it has to process to do all the compression and echo cancellation and speech processing. Therefore they moved all the control and signaling functions from the media gateway to the media gateway controller. Consequently you need to have a protocol between this media gateway controller and the media gateway that allows the controller to tell the gateway to do things like "use this codec and send the compressed audio from circuit 3 on trunk 5 to this IP address."

COOK Report: And what are companies like Lucent and Nortel that are going to use and incorporate this protocol going to do with it?

Rosenberg: They're putting these things into soft switches. Soft switch is a word that bandied about a lot, but it's generally synonymous with media gateway controller or call agent. They all refer to the same box that fits in the IP cloud, talks Megaco to the media gateways and the signaling gateways and effectively controls the media gateways using this protocol. The soft switch becomes a way to decompose the SS7 gateways into a more logical signaling device.

COOK Report: If somebody is doing voice over IP or Internet telephony in the Internet, on TCP/IP networks, is a soft switch essentially, then, a translation device that allows this same person or company that's using

the soft switch to do IP telephony to the entire public switched or global switched telephone network?

Rosenberg: Well, you don't need a soft switch to do that. If that is all you want to do, you could use the Internet telephone gateways that were just PC's with cards. You could buy those off the shelf years ago and you could allow Internet telephony end users on the telephone network to have access to that. But what soft switch is allowing us to do is to scale those systems much larger and to allow them to work better when inter operating directly with the SS7of the PSTN.

COOK Report: In other words, one such device in an organization with a lot of people can enable a lot of phone calls? Does this bring up immediate Quality of Service issues?

Rosenberg: Quality of Service is an issue, but this gateway decomposition does not have an effect on the issue.

COOK Report: So soft switches are useful for a large organization — for example a carrier or for someone who maybe selling Internet telephony or is doing a huge amount of Internet telephony and communication with the PSTN. But they would not be particularly useful for one organization that just wants to enable its headquarters office to do voice over IP to 150 branch offices around the world?

Rosenberg: Yes — exactly. A small enterprise doesn't need to have a soft switch, primarily because the soft switch is really geared for handling an SS7 interface.

COOK Report: So the soft switch is good, then, for a Level 3 or a Qwest, that wants to sell a lot of voice over IP to the public?

Rosenberg: Right.

COOK Report: I am told that ENUM is "hot." Can you take me through what happened with the ENUM working group. And can you help me understand where that's going?

Rosenberg: With ENUM the idea is to "map" phone numbers from the PSTN to Internet connected devices. In the case of me just picking up my phone and calling you on your phone with the call routed long distance over the Internet, we might have a soft switch there that's determining the call, but it doesn't need ENUM because the number that I dialed is already a PSTN phone number and it just pretty much assumes that.

However, if you have a call to a phone number, but that phone number is actually a device on the Internet, not on the PSTN, how do you contact this particular phone number? This is what ENUM is for. One of the reasons it was conceived is that IP phones were going to be given normal telephone numbers.

COOK Report: Do you mean, some numeric string @voip or whatever they would call the top level domain?

Rosenberg: No, it was going to be normal phone numbers that would be doled out to organizations. And that's why you need a database key to figure out, given a particular phone number, what provider owns that phone number so it could get it provisioned to resolve to a particular host or a SIP server.

COOK Report: Oh, instead of area code 44 for England or 7 for Russia, it might be 999 for Internet.

Rosenberg: Exactly. Actually, I'm not sure that that proposal is accepted or it might still be under consideration, but that's sort of irrelevant. In general, if you want to have a user from a phone call a user on a PC, somewhere you need some kind of directory service. That's the initial application. I think perhaps part of the reason it began to pick up steam is people realized that there were more applications than just that one. And of particular interest is the existence local number portability databases. It's fundamentally a database transaction.

COOK Report: So are you saying that in your web browser, Internet-aware telephone, you could have some segment of this database stored and that if you put in some phone number, it would have the appropriate intelligence to connect, to signal, and to get you where you wanted to go?

Rosenberg: The database wouldn't reside in my PC. Instead, like DNS, it resolves and sits out there inside of the network.

COOK Report: Client server?

Rosenberg: Right, DNS is client server. So what ENUM would enable, for example, if I were, from a normal phone or even from my PC, to dial some number which turns out to have been ported to a different provider, then, from the Internet side, the ENUM protocol could do a query, figure out the number, the IP number of where it was ported to, and when I make the call, take me directly there.

COOK Report: Great. How then would you explain what we've just talked about and anything else that's significant that's happened in the last six months or so leading to the kind of meeting that apparently happened in Geneva in January 2000 between IETF and ITU people.

I mean, the message that I've heard is that all of a sudden, the ITU, the European PTT's have got religion and they're ready to go. Is that your perspective or do you have a different one?

Rosenberg: In my opinion, there was no kind of meteor strike or huge event that happened six months ago. Rather there's been gradually increasing meetings and discussions between all parties.

IETF and ITU Meet in Geneva on Convergence Issues

COOK Report: So a reaching of a point that led up to that January session in Geneva has just gradually built up?

Rosenberg: I think it's gradual, but I also think that has now become very clear that both the Internet and the IETF and IP protocols cannot be ignored. They are going to be making a strong contribution to the protocols and architectures going forward. Of course on the other hand, the ITU does know alot about the telephone network and it would be useful to have some input from them so that we do not repeat the mistakes of the past.

COOK Report: So to the extent that the ITU people have had their telephony defenses up against those Internet folk, those defenses now pretty well have been lowered. They've knocked on the door and said let us come into the same room as you and let us play in the same game.

Rosenberg: Well, I think if you look at the NASDAQ share prices, you get the point pretty strongly.

COOK Report: Right.

Rosenberg: The Internet is here and the Internet telephony thing is becoming critical. But I wasn't in Geneva in January and I only saw meeting notes from the meeting and afterwards.

COOK Report: And essentially you've told me that in a general sense, it's probably fairly clear to everyone who saw what you did that the ITU in effect invited a bunch of people from the IETF and said let's see how we can cooperate.

Rosenberg: Yes, and as I said, it's a gradual thing. At the IETF, there have been presentations from ITU on structure of ITU and discussing other collaborative efforts.

Certainly there's been a lot of cooperation and work with each other. RTP, which does media transfer on the Internet has encapsulated speech codecs many of which were developed by the ITU. So there's been co-

operation of ITU people helping IETF figure out how to encapsulate codecs. Such cooperation has been going on for quite some time.

COOK Report: Sure. So if one makes the assumption, then, that the signs seem to be that not only the newer, greenfield companies, but also the older, more legacy oriented, PTT-oriented, ITU telephone companies of the world are now realizing that we need to do voice over IP, that Internet telephony is here. In other words, it's convergence time now. What does that mean?

Rosenberg: Convergence is an interesting thing. A lot of the activity in the past has really focused on convergence at the network layer, if you will. I mean that only in an abstract way.

So convergence means gaining the ability to access to telephone, SS7 systems and similar components from the Internet. A porting effort, like Megaco is quite a good example of this. Making access to SS7 over the Internet a do-able thing. And there have been other examples of that kind of thing.

So that's convergence of the network layer. Just getting telephone service, worked together into the Internet by porting protocols and allowing gateway systems and stuff like that.

Applications Converge with Internet Functionality

What's happening now, which is I think a much, much more interesting notion of convergence than network-level convergence, is the convergence of applications. We are finding that telephony services can be greatly enhanced by combining them and converging them with Internet applications that already exist.

Not everyone has yet realized the importance of Internet telephony, but such realization is coming. And these trends are all pointing to the fact that there has to be an increase in value for the consumers, otherwise why they would purchase the service?

COOK Report: Well, then, what is there to say additionally? Are there only two parts to the convergence issue?

Rosenberg: Well, those two are huge. Internet telephony itself we're only just beginning to understand. And examples of the kind of things that would characterize this network layer convergence, things like ENUM and Megaco — and another network protocol called SCTP.

COOK Report: And the SCTP protocol does?

Rosenberg: It stands for Stream Control Transmission Protocol.

It's a protocol developed by the Sigtran Working Group. Remember I mentioned this SS7 gateway decomposition has three pieces. There's the media gateway, the signaling gateway and the media gateway controller. Megaco runs between the media gateway controller and the media gateway. From the signaling gateway, from the telephone network side, it gets the SS7 messages and it more or less has to just do all the call control.

They needed a protocol to tunnel SS7 messages in order to get them from the signaling gateway to the media gateway controller.. And the Sigtran group was chartered to do that and so they developed this SCTP protocol. It's a transport protocol. The phone network and the Internet touch at the periphery of each. You have telephone gateways that people made which just had analog line cards in them and Ethernet cards in them. They would terminate or act as end systems on both the Internet side and the telephone network side. The telephone network didn't even know they were anything more than end user making a phone call.

So, gradually, what we're seeing that, instead of these things touching just on the periphery of a network, network convergence means that the internal guts of the telephone network and of the Internet are being exposed to each other. And that's what's happening here with this stuff. We now want the Internet to know about SS7 signaling.

And for ENUM, for example, we now want the telephone networks LNP databases and all that to be accessible or reachable through the Internet side, via ENUM. And another one is PINT. Now the purpose of PINT is to enable a few services that allow Internet hosts to actually have direct access to services on the telephone network side. It's best explained by example as to what it does.

One of the services is primarily targeted for something called click-to-dial. In this case assume you're on a web page of some e-commerce and you actually want to speak to a customer service rep without using VOIP. Strictly on the telephone network. You, click on a button on the web site and your regular phone will ring, and when you pick it up, on the other end is the company's customer service rep. You have just made a normal phone call. Except that you call has been launched by a third party control mechanism initiated from the Internet. Now the way that this works on the telephone network is that you have service nodes and SCPs that are able to do this sort of thing and initiate calls. PINT is the protocol that allows an Internet-connected host to have access to some of the controls on the telephony side. This is another case where, there's this function of the telephone network and now we're exposing it into the Internet side.

Now, they're also doing the reverse in a working group called "Spirits." When a phone call gets set up, using intelligent network capabilities, and an SCP gets a notification that the call attempt is made, it all runs right now on local logic. But they want to have it be able to launch queries into the Internet to find out what to do.

And one of the main applications of this is a service where, I'm on the phone because I'm connected to the Internet. Let's assume I'm browsing the web or whatever and when somebody calls me, they normally get a busy signal. So instead what they do with this "Spirits" capability is tell your screen that you have a call and ask, what would you like to do? And the user through the Internet can say, Hang up, transfer to voicemail, whatever.

COOK Report: So this might actually put up a menu or a note on the user's screen with "X" number of choices for the user to click to indicate a decision.

Rosenberg: Exactly. For example, they can continue to tell the switch to connect to the telephony gateway so that the call completes over the Internet. And by being connected to the Internet, now they don't have to miss their phone calls anymore.

But if you think about it from a convergence aspect, it's once again taking the controls that have been in existence in the telephone network for a while and really just exposing them as they are to the Internet side of the house. So that's what a lot of these efforts are about at the network level.

COOK Report: That's very helpful. But how do these network levels fit, with the pretty clearly demonstrated economics of running TCP/IP over glass. The new Level 3, Williams type of optical networks that, in effect, if you're talking about moving bits — and voice increasingly is predicted to become just bits that are moved around on the network — that you better look at your old telephone network infrastructure and you better figure out how to transition to this newer infrastructure that, in terms of the number of bits per dollar that it can move compared to the old one is just orders of magnitude cheaper. Any observations on that?

Rosenberg: Well, yes, but the kind of things we're talking about at the convergence of the application layer have nothing to do with transport per se. They're logic and systems that provide services and call control and

things like this. And, from the point of view of the Internet when you port them over, yes, the voice transport becomes cheaper. But, whatever costs were associated with providing these services on the telephone network, it's not clear that they're going to fundamentally be cheaper on the Internet, because the logic and the software and the call control necessary to implement them, are all still the same.

COOK Report: Well, are not these kind of value-added services that are going to make the telephone companies pretty happy, because they don't want to be selling ridiculously cheap 100 pound bags of salt?

Rosenberg: True, but a lot of the services that our people have been able to enable by a lot of convergence done at the network layer fall into the category of porting these services. Just making them available on the IP side. So I agree with you completely that dealing with the service provider convergence only goes just a little bit of the way. It tends to be about just generally making what exists available on the Internet and perhaps a little bit of some variations on those kinds of things.

The true value add for service providers and for consumers is going to be about new applications and services that are enabled as a result of Internet telephony. That's where we get to this next layer of convergence of application. And that's the space that dynamicsoft fits into.

COOK Report: By all means, let's go into there a bit. But first help me understand things, from the point of view of QOS and the advances in QOS that need to be achieved in order for example, to do QOS across 'Net boxes and gateways. What are these obstacles and are they central to what you're doing?

Rosenberg: While QoS is clearly a problem for voice, if you think about it, QoS doesn't have to get in the way of deploying Internet telephony. Although it is much more of a problem for telephony over the Internet. Now you wireless phone's voice quality cuts out in way that would be generally unacceptable if it were on a normal telephone. But people are willing to live with that. Consequently wireless is widely used, because it provides some enhanced value.

Telephony Over the Internet or Internet Telephony

So VOIP is just telephony over the Internet. QoS is a far bigger problem, because you have to meet the expectations of users for a traditional telephony. That's hard. And while the protocols and architectures for that are beginning to mature, we're still far away from widespread deployment of them to the point where any phone call you make over the Internet would be of really high quality.

However, if you're interested in Internet telephony, where the value add comes from the new features and services such that you may not mind the fact that sometimes the quality isn't as good. Obviously, you want to have QoS, but under these conditions it may no longer be a gating factor for deployment.

COOK Report: But what about latency and delay from a QoS point of view? When I stop talking, am I able to avoid hearing a 2-second long echo of my own words going across the receiver? Is latency fairly minimal now across systems?

Rosenberg: Unfortunately, latency can vary substantially when right now when carried out without QoS. To give you an example of the kind of things I'm talking about, are you familiar with a company called Dialpad.com?

COOK Report: No.

Rosenberg: Dialpad.com is a Internet startup that provides web-to-phone calls that are free. Anywhere in the continental United States. And their business model is they download advertising to you in the client that you watch while making the phone call.

Not only is dialpad's call quality definitely usable, but they have an astronomical number of users and are doing an astronomical number of minutes. What they are doing is changing the way normal telephony works. Yahoo! Is another example. It has voice chat enabled as part of its domestic product. It is widely used, not because the quality is great, but because of the way it works with instant messaging. It presents itself in a way that provides new value to the service.

I can subscribe to you as my friend knowing that when you come Online I will be notified. At that point I will click and make a call to you. So I don't have to call you if you're not there. There's a value added to that.

COOK Report: That question is, with Henry Sinnreich I did well over a year ago, an interesting interview that gave him a chance to excoriate H323 and sing the virtues of SIP and so on. And he drew a block diagram for me. And the block diagram that showed the traditional legacy phone company with the legacy phone system plastered on with an Internet overlay and also with an H.323 overlay. Had something like 14 boxes in it. Fourteen different systems that have to be coordinated and interfaced with each other and what he said, essentially, three different independent networks.

On the other side, what we're going to, he had a much simpler system with five boxes, that basically everything was being done over TCP/IP and Internet and so on, and it was pretty starkly apparent some of the economic impacts that was likely to make.

Can you help me understand what has happened in the meantime at the network provider level for someone — like an MCI? According to Sinnreich they have a 14 box system or they certainly had one. He gave me the impression that they want to move toward the 5 box one. To your knowledge to what extent are any of the major carriers really making progress on modernizing and simplifying their networks in this fashion?

SIP and Convergence

Rosenberg: I guess it all boils down to what momentum behind SIP implementation among the major service providers And the momentum is huge, I'll tell you the truth, nothing short of that. Nearly every major telco that I know of has basically said that they want to be using SIP in their network. From a vendor perspective, the adoption is also tremendous. The best way to judge this is to look at the bake-offs that have been happening over the last 9 months to a year.

COOK Report: And how is the implementation of SIP making this simplification of these more modern networks possible.

Rosenberg: Well, SIP helps on a number of fronts. SIP is architecturally very well-engineered to be just a client server request response protocol. And that's very simple and leveraging systems and services on top of that has been straightforward as a result of the cleanness of the architecture.

COOK Report: So it's the glue or the mortar that enables the various system bricks to be put into a cleaner, less complex, better integrated pile?

Rosenberg: I think that's a good way to look at it. Its broad applicability, its clean, transactional client server model means that the same systems, the same protocols

COOK Report: Go ahead.

Rosenberg: Because as we've discussed, the value of just cheap phone calls is gradually disappearing and there has to be some kind of enhanced value and new services capabilities enabled, otherwise, what's the use?

COOK Report: Because everybody wants to do more than sell a cheap commodity?

Rosenberg: Well, it's not just, I mean, a step up for big transport. You need to do more than just sell the normal telephone service over the Internet, otherwise. We have telephone service. Porting ISDN to the Internet just is not going to sell a lot of services.

Giving Users SIP Enabled Tools to Design New Services

So you need an architecture that can really take advantage of the Internet in terms of creating new value-added services and make it easy to do so. SIP wins on both fronts. The fact that it really has an Internet background, that it leverages http, Mime, and urls and these other cornerstone Internet technologies means that's it's a great protocol for unifying such applications as voice and video on the Internet.

COOK Report: So, in effect, you might say that SIP is evidently is clearly victorious now and at some point in the past six months it probably has become victorious and is that another of the facts then driving the convergence?

Rosenberg: Yes, no doubt. The realization that the convergence has to happen and applications layer as well and that undoubtedly SIP is a tremendously powerful tool for these new applications. Because it borrows from the Web and e-mail legacy, there's a lot of service possibilities for combining Web and e-mail and present instant messaging with voice and video.

In fact, I have a theorem to that regard that I put in my portfolio. I call this the Feature Exponentiation Effect. Henry Sinnreich put it on a slide and called it **Rosenberg**'s Law.

But the basic idea is actually very simple. That a set of services and features that you can provide to your consumer increases exponentially with the set of applications you're combining to provide those features. That if you have just voice and video, you have X number of features you can provide. And to tell you the truth, those are pretty much well exhausted in the telephone network, I would say. That's architecture that's been around for a long time and has pretty much every feature that you could really conceive of that's voice and video.

But, when you throw in Web, you all of a sudden have twice as many different ways you can do things. Because every one of those features now has a way to bring Web into its execution. And then when you add e-mail as another application, you're doubling once again. Now, every one of those older voice and video features, it can have just Web, it can have just e-mail, it can have both Web and e-mail..

So you get this exponential effect as a result of this. This Feature Exponentiation Effect means that there's this large space of features and services, both horizontal features and services and vertical features and services (and I can explain a little bit about how we see the differences between those two things) that are enabled through combining Web and e-mail presence with voice. That means that service providers have opportunities for new revenue which weren't there before in a case where the business model did not go beyond just porting voice to the Internet and doing telephony over IP. This situation enables service provider differentiation. It also means that there's an opportunity for third parties to come in and create services and features which are particularly suited to fit their needs. This is another thing we very strongly believe in. We look at the Web and we ask ourselves, why was it so successful at having so much innovation come about so quickly?

And the answer to that, we believe, is because it pushed innovation to the masses. The whole Internet model of the smart device on the edge, and the Rise of the Stupid Network (I'm sure you're very familiar with that), applies now to the applications as well. The ability to develop new applications and services rests with the end users. You can have all these great ideas come about, because they are now so easy to do. Suddenly users no longer have to depend on their telco to wait two years to roll out some new capability. They could just go get a T-1, hook up, run a web server and boom! Right?

We believe that the same kind of services renaissance really that blossomed in the Web, could very well happen with Internet telephony. By allowing for third parties and end users and system administrators and groups and clubs to create their own services with the right tools. Because there are so many different possibilities of what these things might be, there are just huge possibilities for unleashing a services renaissance for Internet telephony. And that's what we're all about here.

ISPs Become Communications ASPs?

COOK Report: And in one sense the legacy telephone companies may have to come to understand that they don't need to feel that this is the end of the world for them, if they can change their technology and their approach, because although the voice minutes may become a lot cheaper, the applications for the use of this cheap communication just explodes. You increase the size of the market.

Rosenberg: Exactly. We believe that the future service provider is going to be basically a communications ASP — one that maybe even allows its customers to create their own customized services. Perhaps there will be very focused communications ASP's that look at a particular market segment. For example, you could develop a set of specialized services specifically for the legal profession. and have a whole array of Internet telephony services that would be well-suited to the needs of lawyers. , so maybe there's a business model out there for an ASP that provides lawyer-based Internet telephony services.

So basically future revenue is going to be in the service possibilities and the opportunities are in the applications phase. So these providers can win if they realize that their futures are found in being very large, highly reliable, communications ASP's. And that doesn't even need to tie them to the network transport themselves. Just like the Web revolution, where a company like Yahoo! didn't have to own it. It only had to get its content out there. Similarly, a service provider doesn't have to own its own switching infrastructure anymore. They could provide telecommunications services matching customers to the right sets of tools and getting connectivity through their applications.

COOK Report: This is what I realized in talking to Equinex some months ago. Equinex understands that the market is growing very outsourced and very specialized and that ISP's increasingly are vertically integrated combinations of horizontal businesses, that where you paste together an e-mail provider and you paste together this kind of service and that kind of service and the ISP becomes a administrator-coordinator of out sourced services.

Rosenberg: You've got it. Growth in any market happens when a big vertical segment is hacked up into horizontal pieces, allowing companies to piece together the horizontal components from different specialized providers.

And e-commerce is a great example of that. Access to the network now involves choosing from Internet access providers. From the content providers. From the people who make the web servers and web farms. From the people who provide the applications like hosted e-mail. From the billing and transaction systems. All these are broken up and there are specialized companies that provide each of them. Consequently would be dot.coms just have to go pick and choose, they don't even need any technological expertise to have an e-commerce business on the Internet.

COOK Report: It opens the whole arena to who can become a dot.com very easily, dra-

matically. That's an example of pushing the possibilities to the edge.

Rosenberg: Exactly. And that hasn't yet happened in the telecommunications area. But we believe that the future is going to involve this kind of segmentation in Internet telephony. Such segmentation will include the breaking of these vertical markets, allowing application service providers to provide specialized, either vertical or horizontal features and services and allowing telcos to piece together transport, QoS and applications and billing and all that, just as we've seen for e-commerce.

Enabling Market Driven Segmentation in Internet Telephony

COOK Report: So is dynamicsoft, then, providing tools for allowing people to do this?

Rosenberg: You've got it. What we are about is we sell a comprehensive service solution, for service providers and ASP's to put together platforms for providing these services.

COOK Report: And what are some of the specific things that they make possible? Give me an example of some specific products and how they're being used.

Rosenberg: We have a user agent product, which is basically a software development toolkit that allows end systems to become SIP enabled in order to have access to these platforms that provide these services. That's one of our products. And we have a Java and C++ versions available.

Then we have a server solution which is a SIP proxy server and SIP location server, which together provide the network service provider components that it needs to get customers connected to these services and applications. And one of the interesting features of our proxy server product is the ability for end users to actually create their own services, using a graphical tool or whatever, and upload them to the network, have them verified by buyer, and immediately instantiated and available.

COOK Report: Meaning what, exactly?

Rosenberg: Meaning that as you click 'send' from your client to send the service up to the network, 100 milliseconds later, this service is now turned on and running.

COOK Report: And available on your servers for your own clients.

Rosenberg: Right. To take a simple example, I could create a caller screening service. It's not that exciting a service, but it's a good illustrative example. I can sit and I can customize this service based on time of day or particular callers, tell it what I want it to do, and upload it to the network. Then, when anyone calls me, making a SIP call and comes to the server that's providing the service to me, the service is executed and appropriately screened or forwarded to my client, depending on what I specified to the service.

COOK Report: So this could become an telephony enhancement of my middle-sized business. Or, what I think you've also been saying, is that this can become a component part of a chest of tools that someone who wants to offer on an outsourced basis Internet telephony services as part of what an Internet service provider business could offer.

Rosenberg: Exactly. One application is just allowing end users to create their own services.

COOK Report: And an end user might be?

Rosenberg: It could be someone like me. It could be the actual consumer, it could be third party service providers. There's a fair amount of flexibility in the types of services that can be enabled. And it's not just strict voice ones. The language we use to describe these services has the capability to allow users to specify that e-mail should be sent on receipt of a specified kind of call for example. So I could create a new e-mail call filtering service, where when someone calls me and I don't want to speak to them, I get an e-mail notifying me of the call sent to my Yahoo! address or something like that. Or the service language also allows you to do things like, when someone calls and you don't want to speak to them, return a Web page in your response. And that Web page will list urls that they can click to send you mail or reach your voice mail or whatever.

COOK Report: Do you have people who are really beginning to do this yet, that you can talk about?

Rosenberg At the Spring 2000 Voice on the Net conference, dynamicsoft announced that it was working with communications ASPs including Estara and I-Link, and that it had signed a memorandum of understanding with Level(3) to collaborate in SIP-enabling Level (3)'s network. dynamicsoft also announced its Application Server, available this summer, that will provide a platform for creation of innovative new services.

COOK Report: Where and how is the message being gotten out to ISP's, to corporations, to various and sundry people in the Internet telephony field? It sounds like what you've got is similar to the revolution a few years ago in the applications enabling people to design Web pages using a GUI interface, resulting in the explosion of Web pages on the Internet. How is the knowledge of what you're doing being spread?

Rosenberg: That message is spreading widely and quickly. And it's not just you and I talking about these applications. I've been to numerous trade shows, both technical and non-technical oriented, where the message is being heard loud and clear, with a lot of service providers getting up there and saying, hey, this whole cheap long distance thing isn't that interesting. We need to be thinking about new applications, things that really take advantage of the Internet and provide some new tools for our customers.

COOK Report: So you guys have got some products that you're actually beginning to sell now that are doing this and, of course, a lot of the recognition and selling of these tools does begin at the trade shows?

Rosenberg: We've been presenting our eConvergence Server Solutions at trade shows for probably about a year now, With VON being the most important one. I should note that last year our eConvergence Server Solutions won Product of the Year awards from *Internet Telephony* and *Communications Solutions* magazines.

I describe them as server solutions that enable the creation and execution of these enhanced applications. When I think toolkit I think about something like FrontPage which is just where I create the thing. The web server is where it executes. We're not so much providing the tool kits used to create these services as we are providing an execution environment that you actually use to run and manage and deploy them.

COOK Report: And the components that are put together to create this execution environment that I can acquire from you or your competitors are what?

Rosenberg: Well, if you look at the product, we have, that's what you need. You need the access side, you need the user agents in order to get your clients connected to the network.

COOK Report: And user agents are basically pieces of software?

Rosenberg: Yes, you can think of the dynamicsoft SIP User Agent as software development kits that enable SIP-enabled end systems. We also offer SIP Proxy Servers and SIP Location Servers, which are available now. These are the platforms that let service providers deliver connectivity and services.

COOK Report: And these are platforms that are a combination of hardware and software?

Rosenberg: Well, we're a software company. Of course, in the end, everything runs off some piece of hardware, which, in this case, can be either a Sun Microsystems Solaris system or a Windows NT platform.

COOK Report: So if you're buying an e-commerce solution from Sun, might you be buying a combined product that involves some stuff from Sun and some stuff from you?

Rosenberg: Yes, it does. The server products run on both Sun Solaris and Windows NT.

COOK Report: So what are the most important things you should be looking at to evaluate the current state of the Internet telephony market in the sense of where the likely successes will be and where the weak points remain.

Rosenberg: The basic thing to look at is that the successes in the future are going to be defined by the Web model, not the telco model. And that the trend toward innovation, growth and value in the Web way for doing things is going to continue. So look for technologies to win which allow for provision of telecommunication services in a way that is Web-enabled.

COOK Report: In other words, getting the ability into the hands of an end user to make the system do what he needs it to.

Rosenberg: Being Web-enabled encompasses a number of different things, which certainly includes the concept of pushing innovation into the hands of the masses. It also means really almost literally taking advantage of the technologies of the Web in whatever ways are possible. One of the important capabilities of SIP is that it can carry web pages and Java applets. And so that enables a whole range of services.

The kinds of technologies that are Internet-style and make it possible for small start-ups to get into the game, are the technologies in that are going to win. In general that means looking to the Web. They're not beginning to be deployed hugely by the traditional telcos. But if you look towards the Web, that's where you're seeing a lot of the initial deployment of VOIP coming from. Yahoo! Voice Chat does a lot of minutes. And substantial, comparable to some of the Internet telephony telco types.

And Yahoo's application is not even discussed that much, but because they're doing something different with it. Look also at dialpad.com, which is doing huge numbers of minutes. So look to these kinds of companies, the Web model, to see the real initial deployment of VOIP, with the telcos, once again, getting sideswiped by these guys who just aren't afraid to go and run with it.

COOK Report: In the meantime, simply go to British Telecom, French Telecom, Deutsche Telekom, AT&T, Sprint, MCI, NTT etc. Once there if you talk with the key decision makers in those companies, would it be reasonable to say that they're well aware that voice over IP, and all the technology things lower down at the network level, are a real tsunami that they cannot avoid, and for which, at the basic foundation level, every carrier is preparing.

Rosenberg: Yes. I think that's becoming clear and I think it's becoming clear to these companies that the statement can be even stronger than that. There's a dawning realization that it's the applications that are going to really win and I think some of these large service providers are beginning to realize that the Web way is definitely the way they have to go.

COOK Report: The Web way is the way that they're going to have to go and the symptom of this realization is?

Rosenberg: The proof is that many of them have realized that SIP is going to be one of the core technologies that enables a Web based way to do call signaling, features and services. That's why I think SIP has seen so much success in the past year.

Some Internet Telephony URLs

Editor's Note: the following are some excellent IP Telephony web based resources. Most are small encyclopedia giving lists of organizations, protocols, glossaries and a general who's who of the people and player in the industry. These were posted to the IETF list on May 8.

http://www.fokus.gmd.de/research/cc/glone/projects/ipt/

http://www.tsufl.edu/williams/Projects/InternetPhone/TSCIS445.htm

http://www.cis.ohio-state.edu/~jain/refs/ref_voip.htm

See also http://www.cs.columbia.edu/sip
http://www.cs.columbia.edu/~hgs/internet

NOTE: We will be in Russia between May 24 and June 10. We planto publish our August issue on or before July 1st.

ICANN Footnotes

Continued from page 15

vote to reject the broad consensus for 6-10 new TLDs.

3. Proposed ownership arrangements undermine competition and diversity

The new registry proposed by NSI would be cooperatively owned by all the existing registrars accredited by ICANN. This means that NSI, which currently controls about 80% of the registrar market and holds a monopoly on the gTLD registry market, will hold a major stake in the new registry and will profit from its success. The NSI proposal does not specify ownership and governance arrangements, but typically ownership shares are based on market share. Given its size and resources, NSI would have significant influence on the proposed new registry's pricing and policies. Working out governance arrangements among over 100 registrars, with new ones being accredited every month, will not be simple, contradicting NSI's claim that their proposal will speed up the introduction of new TLDs. Even if NSI does not dominate ownership of the new registry, neutral observers must be concerned with the spectacle of a Domain Name policy making body that is only able to award resources to its own members.

4. Collusion proposed

In its desire to protect itself from competition, NSI was not satisfied with reducing the number of new registries from ten to only two. It also proposed to run the "back-office services" for the .banc restricted TLD. In other words, one of the two new registries NSI proposed would be none other than NSI itself.

Other significant points:

"Proof of concept" a deceptive ruse

The NSI proposal is based on the false premise that authorizing a new top-level domain registry is a step into unknown territory. The small number is justified as reflecting the need for "proof of concept." But technical experts agree that there are no technical barriers to adding thousands of new names to the root. Operationally, there is nothing new or untested about adding new top-level registries to the Internet root. The TLD .int for international organizations was added a few years ago with no significant problems. The late technologist Jon Postel, who administered the DNS root for more than a decade, drafted a proposal defining procedures for adding 150 new TLDs back in 1996. In 1994 alone, 50 new country code

Continued on page 26

Commoditizing Bandwidth - Part 3

Commoditization of IP Bandwidth

Some Unresolved Technical and Structural Issues

Interview with Noel Chiappa Emphasizes Uncertainty About Ability of Routing System to Cope with Massive Changes

Editor's Note: Noel **Chiappa** is an independent researcher who developed the multiprotocol router in the early 1980's, while at M.I.T. After leaving M.I.T., he worked with a number of companies to bring out internetworking products. He has been a member of the IETF and its predecessors since 1977, and served as the Internet Area Director for Services of the IESG from 1987-1992. Although semi-retired from commercial activity, he continues to do research on problems in the substrate layers of the Internet, particularly the internet layer. Seeking his assessment of the feasibility of the rapid development of a commodity market in bandwidth, we interviewed him on March 31, 2000

COOK Report: Suppose that it were possible to buy commodity bandwidth in DS3 sized chunks on a month-to-month basis. What would this do to the older carriers? MCI Worldcom, Sprint, ATT and the like? Would many of their corporate customers leave and just buy commodity DS3s?

Chiappa: Let me react to that by making an analogy with corporate use of trucking. A lot of firms that had to ship stuff around would have an in-house trucking department. And then later they found out that it was more cost effective to outsource their shipping. That somebody would take their business, put it together with a bunch of other businesses and get an economy of scale necessary for a cheaper independent trucking operation. But it goes back and forth. Sometimes they decide they don't like that, if it's going to be cheaper, that someone else is making money off the business and they could bring that money in-house and they brought trucking back inside the organization.

Commoditization Changes the Business Model

So I think you're going to see exactly the same kind of thing going on here in terms of whether a large company wants to go out and buy bandwidth themselves and build their own network or whether they want to outsource. It's exactly like the trucking situation.

COOK Report: I expect you may be right, but when this possibility is created, it seems to me that it gives this ongoing segmentation of internet services and structures or "horizontalization," for lack of a better word, some added impetus. For example, if some good size users and corporations wanted to buy their bandwidth directly and didn't want to be a customer of MCI or a big upstream provider anymore, then you get some other things like the decentralization of the technical operation of network infrastructure coming into operation. In other words if I'm a corporation and I can buy bandwidth, since I'm not buying it from a carrier, then I also need to look around locally to see if there's Joe's Local Router and Switch Shop that specializes in the corporations in Hudson County, New Jersey.

Chiappa: Right, you can outsource your "NOC" (Network Operations Center) as well. I mean, you might outsource your bandwidth to one entity and then outsource your knock to another.

COOK Report: Yes, and the possibility, though, what may be holding this back a little bit is that so far the idea of big or middle-sized corporations being able to buy a DS3, for example, haven't been able to buy them themselves on any kind of a predictable price basis. And again, risk management with commodities you can do much more favorably in that you can buy it on a month-to-month basis or three months at a time and you don't have to buy an IRU.

Chiappa: You examine your options, you get all sorts of secondary markets developing. Anytime you commoditize something, you get secondary markets to minimize risks, based on paperization of things like options and things like that.

COOK Report: So all those Joe's Router and Switch Shops may have already started and if they haven't, they soon will be. And, again, this horizontalization will be given some further impetus and you get more and more what Einar Stefferud sees as a true Internet being — how does he describe it? What he calls peer to peer Internet working as...

Chiappa: as opposed to ISP supplied to the customer.

COOK Report: Yes, you get a more distributed, less centralized network, which philosophically to me sounds fine.

The Fate of Backbones Under Commodity Bandwidth

Chiappa: Right, the problem with that kind of peer-to -peer internetworking network is do we have the structure, in terms of the protocols and everything else, to make it work? This is where I become somewhat dubious.

COOK Report: Right and the one other question I have in mind is how do we handle backbones? And do we need backbones anymore? If the big carriers who have the big backbones lose their customers, now what's their business model?

Chiappa: The business model that those guys are going to continue to have depends on traffic patterns. Because if you're Company X and 60% of your external traffic is to Company Y, then it makes sense to set up a private link between you and Y. But if .05% of your traffic is to Company A and .05% to Company B and .05% to Company C, all of a sudden at that point maybe it's not worth having your own private links anymore, then you want to depend on someone else's backbone.

COOK Report: Might there be a middle ground where you look at your total network traffic, it's almost like how ISP's decide whom they peering and whom they do transit with. And to the extent that you have 10% or 20% of your traffic with 3 or 4 people, or 30%, you buy some commodity bandwidth to reach them and you remain an upstream customer at a lesser level of packets shipped to somebody that does have an internet wide backbone?

Chiappa: Absolutely. And the other thing is, too, for a site which is basically a retail site and where they're going to have 17 million people coming in from all over the Internet, they more or less have to maintain their commercial backbone links.

But here's the reason why I see something like backbones as inevitable. I love to make

the analogy from the data network to the road network. Let's assume that you're going to drive to California. You're not going to drive from wherever it is you live to California on the back roads. The first thing you're going to do is get on an interstate and you look at a map. Page one is the Interstate map and the atlas just shows all the Interstates. What people generally do is look on their local map to figure out how to get on the Interstate (although they may know that already). Then they look at the interstate and look at only the big roads to figure out approximately where they're going. And then when they get there, when they get off the interstate, they look at the local map for their local area to figure out how to get from the interstate to where they're finally going.

And we do that for a number of reasons. Number one, those roads were engineered to take traffic long distances at high speeds. Which is to some degree what the backbones do. The other thing is that they're designed in such a way that for most of your trip, all you have to think about, the map that you look at for the long haul is very simple, it's just got the Interstates on it. If we had to compute paths between every little company and every other little company by going through 18 other small companies, the routing problems become very difficult. And one of the functions of the backbones is to make life easier in terms of computing paths. So despite an eventual commodization of bandwidth, just for a number of reasons, it makes sense that something like a backbone business will continue to happen.

COOK Report: But in other words, the big carriers who now have big, big backbones and sprout shoot-like semi-big backbones, if you envision these roots going down to the locality, as more local traffic stays truly local the roots may shrink upward to the big backbones and may focus more on the big backbones. It seems to me the big vertical backbone people become a little bit less vertical and undergo some horizontalization and Internetworking of their own.

Chiappa: That could be. The other place I would look to for commoditization analogies is what companies are currently doing with telephone service. And talking about outsourcing, a lot of companies now, they have their own PBXes. But realize that, for a company, telephone service is as critical a resource as Internet service. It's not like garbage pick-up, where if it doesn't happen, you can sort of bumble along without it. But if your telephones go dead, your company is shot between the eyes. It's the same as your computers going dead or your Internet going dead.

So the Internet, telephone are core connectivity elements in the modern corporation. They simply have to work. And you may be prepared to outsource them, but you want to do it in a way that guarantees you control over them and that guarantees that they are going to continue to work.

COOK Report: And you probably want some redundancy in your outsourcing.

Chiappa: Right. So my question is, what are companies currently doing for their phone service? How are they doing that? What lessons can we learn from that in terms of how they're going to operate in this new model?

COOK Report: Do you have some of the answers?

Chiappa: No, I don't, unfortunately. And as far as commoditization goes the local phone service is not a commodity market, because there's still basically a local phone monopoly. But long distance service does seem to have turned into a commodity market, at least for corporate customers.

COOK Report: Well, even for me, prices are falling.

Chiappa: Right, but in the sense that if you're a big corporation, you probably shop around every couple of months to see who's got the best deal on long distance bandwidth. And you may, I don't know how they work it, whether they switch every two months depending on who's cheapest or what.

The third comment I wanted to make about commoditization is that anytime, you look at a market that becomes commoditized in which the resource that is being sold becomes fungible, you get a very different business model. For example, look at computers back in the 60's. You were tied into your vendor's operating system, and because of that you had to go back to him to buy more computers.

Portable operating systems come along and, presto, all of a sudden the computer industry is a commodity market. And look what that did to a lot of these mid-sized sources of computers. They just blew up and disappeared. There are far fewer computer architectures in the world now than there used to be.

COOK Report: But on the other hand, computers are far more ubiquitous than they used to be.

Chiappa: I understand that, but in terms of the number of players in the market, every time you commoditize something, what happens is the number of players in the market goes down and the size of those players goes up. Whether it's airlines or oil companies or computers or whatever. Commoditization means you have fewer, bigger players. Because when you get into a commodity market, the only thing that matters anymore is price.

Look at the computer market. When it comes to PC's, there are a couple of CPU vendors who have a niche in that aspect of the business. Intel has the majority of it; there's Power PC's and a few others that are still nibbling around the edges, but basically it's all Intel.

Now, that's an interesting case in point, because ten years ago, Intel had the PC market and the high end market was held by other people who built specialty computers. Now, technological change has allowed Intel to creep out of their low-end niche and go up and eat the lunch of the people in all the middle place niches. So people like SGI basically got their lunch eaten by Intel.

Chiappa: And some of the real high end guys are still there, but even those, you know, Cray's gone under. Basically, who's left? Other than Motorola and Intel, who's left? IBM. But who else is left? I mean, essentially, the number of niches can diminish, too.

Routing Issues in a Bandwidth Commoditized World

COOK Report: Is there anything that you see, but from the point of view of routing or routing protocols? Are there any magic bullets in routing here that are going to get more important than any other magic bullets in a world like this?

Chiappa: Yes. The one thing that is most likely to slow the commoditization of bandwidth is that the fundamental Internet architecture, you know, the protocols and the routing and all, isn't ready for it yet. It's all spit and bailing wire.

COOK Report: Oh, yes, that's Cook's favorite phrase nowadays. Absolutely.

Chiappa: Right, we've got a system that's held together by such delicate threads, it's hard to take it and start introduce serious changes.

COOK Report: And you've got the IETF and IAB and IESG. I don't know what your opinion is, but I am increasingly hearing that it's so politicized at this point that it can't do much that should be done.

Chiappa: What's right and wrong with the IETF is a whole separate topic. The point I was just going to make is that commoditization implies a tremendous amount of flexibility, a high rate of change

in topology in the way things are connected together. It also requires a very robust infrastructure. That's one thing we certainly do not have. A lot of this stuff runs because there's a lot of smart people tweaking it all the time. And I'm not sure we have enough smart people.

COOK Report: To what extent do you think Cisco could begin to have a business model that, I mean, Cisco is losing out on the routing business to Juniper, are they going to abandon the huge backbone router market to somebody like a Juniper? And my friend Dave Hughes, who's been invited out there and wined and dined says: boy oh, boy, Cisco is now really into wireless in a big way and with lots invested in routing protocols in people and in companies designed to hold the middle together, as opposed to the inter-city interstates. Is there anything to that?

Chiappa: Cisco, being a big company, has to be interested in markets that are big dollars. And you make a lot more money by selling a million things for $1 than you do by selling ten things for $10,000. So I think they probably perceive the mobile market as something with more long-term growth potential in terms of the sheer numbers.

COOK Report: knowing the talent, the incredible intelectual talent that Cisco has hired recently, if there's a problem, as you said, of spit and bailing wire and mid-level architectures and a need to re-architect routing to enable routing to work in the commoditized bandwidth market. If Cisco is as smart as we would think they would be, wouldn't they have a bunch of people working on how to develop routing that would make their diverse family of products click with each other in jillions of places around the world?

Chiappa: Here's the standard spiel I've given people over the years. I don't know that Cisco is doing any such thing. Doing better routing, it's not a technical problem, it's a political problem. We know how to do it. We've known for years what the outline of a much better routing (path selection) architecture looks like. The problem is getting the IETF community together to do it. Because it's like deploying a better http. If you get only 5% of the world to deploy better http, it doesn't do you much good.

COOK Report: Well, if they rolled something out and they did a good job of it and they say any Cisco person who buys our products can route in this fashion, then don't they have a proto-standard where people will begin to hop on the bandwagon?

Chiappa: Yeah, but the hard part is not getting to 5% and the hard part's not getting from 90% to 100%. The hard part's getting from 5% to 70% or somewhere in there.

Can Cisco Solve the Problem?

COOK Report: But Cisco has so many products in such a wide, diversified installed base of such a large family of products around the world, that if somehow — and there may be some reason it would extraordinarily difficult to do that. But imagine if Cisco threw out their entire family of TCP/IP conversant products, could it develop an inter-family, inter-router wireless and fiber compatible routing mechanism that wasn't held together by spit and bailing wire? And do it for everybody?

Chiappa: Well, here's the problem. The problem is, they've still got to interact with everybody else.

COOK Report: Well, you don't have end-to-end connectivity, then. But you do have the possibility of the big Cisco part of the world saying, well, Cisco then has the rough equivalent of a NAT box. Not only do we have stuff that connects our stuff together, but we have special boxes that translate our stuff to the other stuff still that's out there in the world.

Chiappa: Yes, but here's the problem. Let's say I want to go from my company to some other company. And I've got to go through a Juniper box to get there. The places where you're going to be selling bandwidth is not on the edges. The places where you're going to be selling bandwidth to some degree is in the core.

COOK Report: Yes, you're right in one sense, but with the Akamaization of the world, how much of the total bandwidth is out on the periphery and hops from one periphery point to another without going through the big core? Now, your big core is going to remain huge and that's true. But it's going to be somewhat flatter, I bet, than it is now.

Chiappa: Gordon, I've been thinking for the last couple of years about what's the business model that works for the deployment of a new routing infrastructure. If I had a simple answer, I would have done it already.

Here's the problem. When the Web first came through, it had what I call self-deployability. Okay, so what does self-deployability mean? I'll tell you what it means. Imagine the world when nobody had the Web. One percent of the user base deploys the Web and has browsers on their machines. They get immediate benefit from having that stuff on their machines and using it. Even though 99% of the rest of the world doesn't have it, the one per-cent still gets a great benefit from it. So immediately it sort of proselytizes itself, it almost deploys itself. Once 1% has it, then it becomes 2% and 4% and 8% and 16% and, very quickly, everybody has it.

COOK Report: Well, couldn't you make an argument that if Cisco rolled out a good proprietary routing architecture for its products, no?

Chiappa: That's the problem with routing. If 1% of the world deploys a new routing architecture, it doesn't do them very much good at all. So how do you get to 2%?

COOK Report: I don't now.

Chiappa: If you could solve that, you can tell me how we can deploy new routing architecture.

COOK Report: Is there any reason why Cisco couldn't perhaps develop a new architecture to unify all its products?

Chiappa: Yeah, but the problem with the routing is not making it run on the small scale. The problem with routing is not making it work to your neighbor. The problem with the routing only comes in when you try and make the entire system work.

There's no traffic lights at the end of my street. If I want to drive, I can drive through my neighborhood, I can get places at 5 o'clock on a Friday by taking side streets where there is no congestion. If I take main roads, there's congestion. That's a problem that's a system problem, not a local problem.

COOK Report: So even if you have Cisco, "cities" all over, where all Cisco protocol architecture routing boxes talk happily to each other, if you assume that the cities have to talk to each other, then you've got a problem of how you do that? And if you assume there's going to be this magical NAT box at the exit or entry gateway to each city, is the imponderable problem how in the hell do you get stuff out of that city architecture, through that magic gateway an inter-city type of thing?

Chiappa: But the problem is slightly worse than that. The problem is, let's assume you have all these little cities running Cisco software. And what I was trying to tell you is that you're never going to get to that stage and here's why. Because to get stuff around locally, the stuff we've got works fine. You don't need to deploy Cisco's new software to get the stuff we've got around locally. Nobody's going to bother to deploy the stuff to do local traffic, because local already works fine. You only need this new stuff for long haul. But it won't work until everybody's got it. Catch-22.

COOK Report: Yeah. Well, it's like why

IPV6 isn't going to get anywhere, unless somebody becomes Czar of the Internet and commands it and how are they going to do that?

Chiappa: I doubt that anybody is working on anything and I think the problem is twofold. Number one, nobody sees the business opportunity in it. And number two, it has this problem that I explained that it's like the Catch-22, it's no good to everybody until it's deployed everywhere, but it won't be deployed anywhere, because it's not very good when it's deployed locally.

Is Infrastructure Sufficiently Robust?

COOK Report: But you're also implying that there's going to be limits into which the commoditization of bandwidth eventually is going to run.

Chiappa: Right, because the infrastructure is so crummy. The commoditization of a lot of products depends on a lot of infrastructure. The commoditization of oil depended on good pipe lines.

COOK Report: Well, we've got a lot of infrastructure, but you're saying it depends on very smooth and very reliable infrastructure, because your point is, if I spend $10,000 for an OC3 from point A to point B for tomorrow evening from 10 to midnight, goddammit, I want the infrastructure to ensure that it's going to work, don't I?

Chiappa: Well, not only that. You want the routing to adjust very quickly so that you can use it and when it goes away, you want the routing to adjust back.

And that's my question: can the routing really do it? The sense that I have watching what happens in the Internet when there's an outage now, because I'm sitting down here in Virginia. And I go over my local ISP, which goes to a major backbone, which sends up to another major backbone, which then jumps to MIT. And I watch it when something breaks. And the answer is, that it looks really ugly what happens. And I understand why some of that ugliness is happening. And this is why you need to talk to people like Michael O'Dell and Sean Doran and try to understand more why the Internet routing is set up the way it is today. It just reacts really slowly and really poorly. I mean, this stuff all ought to...

COOK Report: Yeah, Sean tried to describe his sort of hierarchical philosophy and Mike O'Dell's ATM mesh philosophy.

Chiappa: No. Don't ask these guys about how the routing works inside their system. You need to understand from them how it's working as an overall system, how it works between ISPs. That's what's driving some of how it's working inside. In an attempt to reduce the routing overhead load caused by inter-ISP route flaps, the dampen the network of the offender. Which is to say that for a specified period of time they refuse to listen to any announcements. Because it's so painful to get hit by a dampening filter, people have turned the knobs on their IGPs down, so just to make absolutely sure that no change is reported to the inter-AS routing, unless there really, really, really, really, really, absolutely is a problem locally that can't be controlled. And the problem is, that this behavior has all contributed to a horrible response time when something really does break.

COOK Report: Because if you've got a break, you've got a bubble in your system and you don't want that bubble to hit a system border and spill over into somebody else's system.

Chiappa: Right, unless it's really a major problem that you have to tell them about. So they're very conservative about how quickly they tell people they've got a problem.

COOK Report: And BGP is what controls the intersection of UUNET's border with Sprint's border, while IGP is the protocol that selects the path across the inside of the network?

Chiappa: Yes. "I" stands for interior. There's a number of different IGPs. RIP is one, IS-IS is one, OSPF these are all different IGPs. Now, there's another class of protocol which is used to talk between autonomous systems called exterior gateway protocols (EGP's) and at the moment there's only one and it's BGP. So what I see happening is, I look at the routing when something breaks. I look at how long it takes for everything to recover. And it does not give me a warm and fuzzy feeling about how robust the routing is on a system-wide basis.

Problems of Turning Big Pipes On and Off

And, to me, that's going to be one of the critical gating factors on whether or not you can commoditize bandwidth. If you start dorking around with the topology of the network, are the protocols and everything else in it robust enough to deal with that kind of sort of brownian motion in the connectivity?

COOK Report: Can I make an assumption that if I want to commoditize bandwidth to provide a broadcast or an auction or videoconference or teleconference or something between two cities that I might tunnel that bandwidth effectively and not have to worry about how the whole Internet system accommodates to it? But then do we still have the Quality of Service problem that if I'm all on UUNET's network, I can tunnel okay if I don't have to cross into somebody else's network. What happens if I'm trying to tunnel something and I have to go into somebody else's network, I've got a problem, yes?

Chiappa: Well, yes, but if key buyers of commoditization would be end users who wanted to talk among themselves directly. The routing isn't as much of an issue there, that would probably work. Yes.

COOK Report: Right, and if they want to talk to themselves directly, how are they going to do it? They can go to somebody like a Williams who can sell them a million different colors of light and they can buy a wavelength at the basic physical layer of a commodity fiber provider.

Chiappa: Right. This goes back to something we were talking about earlier, which is expertise. Do they have the in-house expertise to take that link up and down and tweak their routing to make that traffic make that link? Or do they have to buy it from an outside supplier who's has the ability to install that kind of link and have it come up and come down and have it affect their routing correctly? When you install a link like this, the routing has to adapt before the packets will flow down it. The routing is what's controlling where the packets are going. Until it's finished adapting, no packets are going to travel down this link to this bandwidth that you've just installed.

So I think commoditization of bandwidth is a good idea. How quickly it will affect things is going to be gated by a number of factors, one of which is how good, is the network structure, are the protocols and everything else really ready to deal with that kind of network? Are they robust enough and flexible enough to handle that kind of network? The answer is, I'm not at all sure they are.

ICANN FOOTNOTES

Continued from page 22

TLDs were added to the root, and country code TLDs operate in the same way as .com/net/org TLDs. There is simply no evidence for NSI's claim that new TLDs require "proof of concept."

The European gambit

NSI's proposal to locate the new .shop registry in Europe is a calculated attempt to win political favor for the proposal. US-European rivalry has played a major role in the domain name wars, and the Europeans are extremely sensitive to their status. However,

Continued on page 31

Deployability of IPv6 Debated on IETF List

Obstacles Are Huge Cost, Lack of Benefit from Partial Deployment

Complete Convergence of Voice & Data Seen as Impossible Without IPv6

Editor's Note: On the IETF list the IPv6 versus IPv4 and NATed-end-points religious wars continued.

On April 25 **Brian Lloyd** asked: whatever happened to IPv6? 128 bit addresses would certainly allow us to continue using IP addresses as endpoint identifiers thus eliminating the need for NAT. It seems that this is a more reasonable solution than trying to make NAT work under all circumstances.

Noel Chiappa: The basic key *architectural* problem with NAT (as opposed to all the mechanical problems like encrypted checksums, etc, some of which can be solved with variant mechanisms like RSIP), as made clear by Keith's comments, is that when you have a small number of external addresses being shared by a larger number of hosts behind some sort of "address-sharing" device, there's no permanent association between an address and a host. It's *that* that causes many of the worst problems - problems for which there *is* no good workaround (because the problem is fundamental in nature).

Now, if you have a site which has more hosts than it can get external IPv4 addresses for, then as long as there are considerable numbers of IPv4 hosts a site needs to interoperate with, *deploying IPv6 internally to the site does the site basically no good at all*. Why?

Because for interactions with those external IPv4 hosts (who will be the vast majority of the hosts one wants to talk to, in the initial stages of deployment), *you have exactly the same architectural problem*. No matter what IPv6<->IPv4 interoperability mechanism you use, you still have that same *fundamental* problem - no permanent association between a host and an address (in this case, the IPv4 address that it *has* to use to communicate with an IPv4-only host).

When one looks at the overall business/economic case for deploying IPv6, in the light of this, the results are fairly devastating - and explain perfectly what we've been seeing for the last couple of years (rapid increase in the number of NAT boxes, and basically no traction for IPv6).

A site considering deploying IPv6 is in one of two cases: it already has enough IPv4 addresses, or it doesn't. In the foremer case, what's the upside to deploying IPv6? Autoconfiguration, etc aren't enough to outweigh all the costs of switching (to software which is less available, less tested, less tuned, etc).

In the latter case, it's equally as bad: they are going to have to struggle with the problems inherent in IPv4-address-sharing technology whether they go with IPv6 or not, and again, the remaining advantages of IPv6 (autoconfig, etc) are outweighed by the costs.

I'm still sorting through the implications from this, trying to put them all with equal clarity, but one thing that does seem clear is that this kind of upgrade model is economically unworkable in the current large-scale Internet. Exactly what will work is something that needs to be pondered for a while.

One possible lesson is that we need think about how any new stuff is going to make peoples lives significantly easier overall as soon as they start to deploy it, because without that, probably very little is going to get done.

Matt Holdredge (Lucent): we've been through all this already ... at the IAB Network Layer Workshop. One of the conclusions is that an IPv6 network NAT'ed to the IPv4 Internet isn't any better than what we have today with IPv4-NAT-IPv4

Chiappa: Well, my statement is broader than that. It says that *any* IPv6<->IPv4 interoperability mechanism is going to have the same fundamental problems as IPv4<->IPv4 NAT. I think that's a pretty powerful statement, one that puts a hard ceiling on what one can hope to accomplish (in any moderate timeframe) with *any* alternative to IPv4<->IPv4 NAT (including IPv4 RSIP).

Holdridge: So if you are NAT'd to the public Internet today, you shouldn't have a problem with converting internally to IPv6. At least from an architectural sense. :)

Chiappa: Sure, you're going to have basically the same service externally, if you are using IPv6 internally, as you are if you are using IPv4 internally. So, you're the CIO for Foondoggle Corp, and you're trying to figure out whether to spend any of your Q3 funds on IPv6 conversion. Let's see, benefits are not very many (autoconfig may be the best one), and the cost is substantial. OK, let's put it off till the next quarter. Go back to step 1.

IPv6's claimed big advantage - a bigger address space - turns out not to be an advantage at all - at least in any stage much short of completely deployment. IPv6 deployment is going to have to be driven by IPv6's *other* features, and when you take bigger addresses out of the cost/benefit ration, I'm even more dubious that the features that are left (autoconfiguration, etc) outweigh all the costs and risks of IPv6 conversion.

It seems that you can postulate whatever level of IPv6 deployment you like (a long stretch in itself, but just for the sake of argument, let's make it) - 5%, 10%, whatever - and there's still no mechanism to drive further deployment.

On May 7, 2000 on IETF list **Keith Moore** wrote: for a long time the assumption was that IPv6 would be deployed first in the core, and then in the periphery, of the net. I'm now of the opinion that IPv6 will be deployed first in the periphery - both in emerging networks that need large amounts of address space, and in existing IPv4 nets using 6 to 4 - and it will be deployed by folks who have applications that need global address space (and which perhaps aren't already widely deployed using v4) and by folks who need to be able to access the new IPv6-only networks. the emerging networks may be large networks in parts of the world that are just now getting on the Internet, wireless networks, and other networks designed to support large-scale data gathering. (power meters, auto traffic monitors, environmental monitoring, security systems, etc.)

I think we will have a long period of v4/v6 coexistence, with v4 becoming more and more NATted and popular applications moving over to v6 based on how poorly they work under NATted IPv4. the older and better established the application under IPv4, the longer it will take to move it to v6. SMTP will use IPv4 for a very long time - not that it won't use IPv6 when available, but for a long time you'll need to have at least one IPv4-based SMTP server acting as a mail exchanger for your domain, in order to reliably receive mail.

The core will support v6 when it makes economic sense - i.e. when top tier ISPs can save enough on bandwidth and support costs (as compared to tunneling) to make the investment worthwhile. which is not to say that some major ISPs won't support IPv6 before then.

Sean Doran responded: Perry Metzger had this to say a long time ago (1999 12 03):

"Peter made the absurd statement at DC that he'd be willing to provide v6 at some high multiple of the price of v4. Why should we bother? I can just pay 5% more for the extra bandwidth encapsulation will consume and ignore you until such time as you decide it is in your interest to offer native service."

Doran: Clearly he agrees with you that the core of the Internet can effectively run IPv4 ever, or at least until there is a clear advantage to running IPv6.

Peter Lothberg, meanwhile, (Doran continued) has proposed a price which would make it worthwhile for certain ISPs to become dual-protocol. I'm sure others would be interested. Maybe you guys can convince the U.S. and European Taxpayers to pay this cost through direct and indirect government grants and subsidies to ISPs and ISPs' customers, sort-of like what used to happen in the OSI days?

Moore: As for your AM vs. FM analogy - there are a variety of theories about this, ranging anywhere from artifically making v4 addresses even more scarce to encouraging a run on v4 address space and making them scarce that way. but I think the shortage of IPv4 address space will encourage adoption of IPv6 even without changing allocation policy.

Doran: I would like to see a market develop for IPv4 addresses, along the lines of the late PIARA work. This would also encourage a market for routing-table entries, both of which would produce a significant incentive to dramatically improve upon on-the-fly host-renumbering.

There is no reason to believe a PIARA-style market for IPv6 addresses and routing-table entries could not also be interesting and perhaps useful.

There is clearly a "price" associated with receiving a TLA allocation, namely the compliance with a number of IETF-produced rules with respect to how one conducts one's business. I counterbid $1000 in U.S. currency. Sean.

P.S. By "routing-table entries", I mean of course, not just the consumption of memory and CPU resources in forwarding packets in to large numbers of possible destinations, but also the cost in various resources (bandwidth, CPU, complexity) of acquiring and propagating information which may lead to routing-table changes.

On May 8, **Bill Manning** wrote: There is near zero value in the number/address and very real value in the routing slot. Perhaps it is best to simply have ebone route filter on the /16 boundaries to drive home your point. (being cranky this morning)

Doran: I utterly reject your ostrich-like position on this matter. I would be extremely happy if I could make a money-based or (better still) capacity-based offer to one of the R&E networks or institutions which retain very short prefixes (historically known as Class A and Class B networks), without being prevented from engaging in such a private transaction by the collusive behaviour of IANA and the registries.

Likewise, I think that government agencies in various East and South Asian countries, and perhaps various Asian ISPs or next-generation mobile telephony organizations would be extremely happy to bid for a few historical Class As now being under-used by the current registrants, rather than being told "no" or "first demonstrate usage" by APNIC.

It has been made clear in the past that any transfer of addresses will be reviewed by a registry, and that if the ultimate recipient of transferred address space wants more address space from the registry, they must comply with the ordinary "growth & design" rules.

I hear anecdotally that the threat of a withholding of new allocations to the selling party has also been made in the past.

This is a system which enforces a "one-seller" (the IANA), "one buyer" (one may return addresses to IANA only) model, which flies in the face of free markets, and perversely imposes costs upon consumers.

Although I am happy that there are people trying to conserve IPv4 addresses and also encourage sensible routing announcements by providing not less than a sizeable aggregatable range to qualified buyers, the qualification process is tricky and gets trickier as one's business grows.

There is a VERY real cost — most notably in terms of time — to using the "growth & design" scheme for acquiring more addresses than an initial allocation. This has, in fact, slowed the deployment of independently-routed subnets owned and operated by a single organization. This slowness could have been avoided if a market for IP addresses existed, and I can assure you that a fairly sizable amount of money would have been spent to speed up the process of acquiring a handful of relatively long prefixes. I am also aware of anecdotal reports of organizations who had to suspend turning up newly-acquired customers because they could not quickly acquire new addresses from the monopoly vendors: the local registries and IANA.

In my opinion, following the PIARA work, the appropriate thing for IANA to do is to spin off its IPv4 address allocation function. We will call this spin-off IANA-I. IANA-I should then proceed auction off the *ENTIRE* _not-yet-allocated_ address space, being very clear that what is sold at auction is merely an exclusive registration of a range of IP addresses in an IANA-maintained and publically-accessible document, the right to make future changes to the registration, and the right to transfer future registration change rights to another party.

The IANA-I or its agents could certainly charge a small fee for processing such changes from persons duly authorized by the registrant the IANA knows, however it should not have the power to refuse any transfers of title.

I would furthermore like to see the ENTIRE unallocated IPv6 global unicast space auctioned off in a similar manner, abandoning the anti-market "one-seller" model put forward in RFC 2450 section 5.0, using the monopoly tariff put forward in section 5.2 ibidem.

That ISPs probably cannot be compelled to consider the IANA-I registration document at all, in whole or in part, when configuring their networks' routing policies, should be declared by the IANA and its auctioneer agent, much as the registries note this now when making allocations under the current "one-seller" rules. Whether there is value in such a risk-bearing instrument,however, MUST be determined by buyers, not by IANA, IANA-I, or ivory-tower academics.

There are certainly MUCH riskier instruments traded regularly as assets on exchange markets throughout the financial world. Moreover, the IPv4 black market that DOES exist, as noted by David Conrad, argues strongly in favour of testing the "white market" in a sensible fashion.

Finally, a small initial registration fee by IANA-I could allow currently allocated address space could be noted in the IANA-I registration document, thus normalizing the "deed" to the range of addresses, likely making it easier to undertake a transfer.

On May 8 **Sean Doran** also added on the inet access list: One of the bits of fallout from the ongoing warfare in the IETF mailling list is this draft:

http://www.ietf.org/internet-drafts/draft-ohta-address-allocation-00.txt

Some of you may like one of the proposals: No IPv4 address space should be allocated to an ISP, unless the ISP support fully operational fully transparent IPv6 service with at least 64K IPv6 subnets to all the end us-

ers.

Ohers may want to reflect upon how the ongoing artificial scarcity of IPv4 addresses is being used to force ISPs to conform to various peoples' political ideals with respect to IPv6 and other things.

Michael Dillon: Since most of us here are not aware of the ongoing warfare in the IETF, perhaps you could give us a rundown on the players and their positions?

The draft you mentioned and this document http://www.real-internet.org/whatisric.txt at the Real Internet Consortium make interesting reading but I'm not entirely sure where it is all going. There are interesting echoes of Vadim Antonov's router architecture and the Ebone's zero packet loss architecture but is it safe to assume that the players line up in two camps, one clustered around simplicity and parallelism, the other clustered around complexity and reserved/managed flow paths?

Avi Freedman writes: no antagonism intended here, just a question.

Doran: Understood.

Freedman: I thought you didn't like the idea of smaller ISPs getting address space on their own?

Doran: There is nothing like returning to a place that remains unchanged to find the ways in which you yourself have altered. - Mandela

There are a couple premises:1/ the only scalable routing system known now is hierarchical routing, which relies upon aggregation 2/ in order for hierarchical routing to work, network location names ("addresses"/ "locators") must follow the network topology 3/ non-multihomed ISPs are topologically identical to non-multihomed end-users, and both should be aggregated behind a locator which describes the largest possible piece of the network's topology (more simply: you want to be able to have a default locator; likewise, where you need more than a default locator, you want to be able to use a small number of highly-aggregated locators, as well as more-specifics). In other words, where one can use aggregatable addresses, one should. Where one cannot use aggregatable addresses, one needs addresses. That hasn't changed much.

However, it is getting very hard even for ISPs which are multihomed, and which aren't that small, to get large enough address allocations. It is also very hard (and expensive in terms of time and effort) to acquire larger chunks of address space as one's needs grow.

Freedman: Given that, "artificial scarcity" confuses me.

Doran: There is no way to acquire addresses if you do not meet the requirements of your local registry. There are people who "own" huge chunks of address space (historical Class A networks) who cannot _sell_ you address space, when you cannot comply with the registry requirements.

This is artificial scarcity: one seller (the registry), one price (usage-based).

Freedman: Because while I think that any legitimately multi-homed ISP should get a space reservation, I do think that IPv4 space is less-than-a-decade or maybe even less-than-5-years until crisis.

Doran: Yes, any multihomed ISP should get space; we do not disagree.

I don't agree with your timeline, and neither does the ALE work, however I could accept a decade as a very early date of complete exhaustion. Bear in mind that the work kc [Klaffy] and Hans Werner Braun and company at CAIDA have done following on from the ALE work indicates that only a small fraction of IP address space is actually ROUTED anywhere, from the perspective of the Internet.

There are lots of organizations who own old Class As and Bs which simply are NOT ANNOUNCED to the Internet, and likely never will be, until someone can buy or otherwise acquire those Class As.

The inability to buy those addresses without jepordizing they buyer's and seller's future transactions with the IANA and its registries is an artificial scarcity. (Maybe someone could put this more simply?)

Meanwhile also on May 8 back on the IETF list Sean Doran wrote: Ohta-san with regard to <draft-ohta-address-allocation-00.txt>

While I agree with you that the current usage-based allocation system is wrong, your draft's "Assignment Plan" (not more restricted) proposes to continue an anti-market single-seller model for IP addresses of both IPv4 and IPv6 flavours. There is no scope for negotiating with the monopoly seller, given this tariff.

On the other hand, I do particularly like The More Restricted Assignment Plan: No IPv4 address space should be allocated to an ISP, unless the ISP support fully operational fully transparent IPv6 service with at least 64K IPv6 subnets to all the end users.

Because that will force IANA out of its ostrich position with respect to being a monopoly seller with a non-negotiable monopoly tariff that imposes significant costs upon consumers, by immediately forcing the monopoly to stop "selling" addresses except to people who meet extraordinarily onerous and expensive conditions.

Unfortunately, because you do not actually propose only the More Restricted Assignment Plan, your draft effectively continues the objectionable practice of deliberately introducing artificial scarcity into IPv4 addresses in order to force your politics upon ISPs and other businesses. This is identical to a monopoly which has goods to sell but nevertheless deliberately restricts supply in order to support higher prices.

There are two main differences between your draft's proposal and the current system.

One of the differences is that your political ideals include the deployment of IPv6, which is something the current usage-based allocation system does not. This is simply a change of the monopoly tariff, the "price" at which the monopoly will "sell" consumers (non-transferable) address ranges. So, while it is an important difference, it is not particularly interesting, since it is just a higher price in view of a smaller supply.

The much more interesting difference between your draft and the status quo is that artificial scarcity of IPv4 addresses would evolve as a result of over-allocating IPv4 address space to applicants, rather than attempting to allocate the smallest workable amount of address space, as is the practice now.

This erodes IANA et al.'s monopoly supply, because now there is a surplus held by many other parties, who then can act as alternative suppliers of IPv4 address space.

If steps are taken to avoid the development of a massive black aftermarket for IPv4 addresses overallocated by IANA et al., by providing the mechanisms of a "white market" — notably a public registry of IP address title, with an exclusive but transferable right to transfer title to another party — then I would object much less strenuously to your draft, since it is fundamentally PIARA, but with a rather odd auctioning system for the remaining not-yet-allocated IPv4 address space.

Given the involvement of one of your coauthors in the original PIARA work, I am not at all suprised that the draft can easily be read to favour the ultimate development of a market for IPv4 addresses.

Let's just not make that market completely black, with all post-IANA/registry-allocation transactions completely sub rosa.

Note that the development of a "white market" public registry does not rely upon the

IANA. If the IANA and its registries were to immediately cease offering IPv4 addresse AT ALL, it is quite clear that market forces would arrive upon a suitable solution rather quickly. Given that the initial allocation prices proposed in your draft are extremely onerous when combined with the conditions in RFC 2450 ("higher cost of acquiring bundled goods, rather than only one good individually; cross-subsidy"), I imagine that there could easily develop a situation in which IANA et al. simply could not find a buyer prepared to meet their price / qualified to meet the allocation conditions.

Bear in mind that the IANA IP allocation system has two functions: 1/ prevent namespace collision 2/ provide one of many inputs which network operators may choose to use when configuring their networks

Both of these functions can easily be done elsewhere. There just has been no reason to do that yet. Your draft would supply a very strong reason, therefore I support your draft.

How do we get it adopted quickly, and get the IANA, APNIC, ARIN and RIPE to IMMEDIATELY cease offering IPv4 address space to people who do not FULLY comply with the requirements in your More Restricted Assignment Plan, and the various RFCs and standards-tract documents it rests upon?

Editor: Meanwhile on Nanog on May 9th, we find **Tony Mumm** writing: Hop-by-hop routing is on its way out....and not soon enough.

Vadim Antonov: Worked fine for the last 20 years. Can you substantiate your assertion? So far all alternatives were shown to bring more problems than improvements.

Mumm: I am certainly not one to forget where we came from. It did, and will, have its place in the network.

However, more end to end knowledge is required in a converged network. I am for one looking to ease the load of general traffic engineering. In the giant carrier backbone, it may be better to use your resources wisely, than to overbuild your network.

I hear a phrase quite often "Everything is going to IP". Well, before my phone call is on IP, I want some guarantee that its getting from point A to point Z at the rate I'm paying for.

Craig Partridge: I hope you won't get upset if I use your short comment here to get on a high horse and rant for a moment.

Over the past several years, I've heard several people say we need to embed more end-to-end knowledge into the network, as a solution to quality of service, or reliability, or some other valuable function. What I have not heard yet is:

1. A cost-benefit tradeoff. Embedding end-to-end cognizant information into the network has a cost, often in reduced flexibility (e.g., look at how hard it is to add new applications to the telephone network vs. the IP network).

2. A reasoned technical justification that shows that we can't provide the same service with the current service model (which I define roughly as "route based on the contents of the IP header") and that we need to break or bend the current model to do new things.

Let me give a concrete example. It is fairly clear that one of the advantages of packet switching is that it allows us to build fairly reliable networks out of much less reliable parts. (Viz: the Internet is getting closer and closer to the reliability of the telephone network, yet no one claims that a Cisco router is as reliable as a #5ESS). Yet, oddly enough, we don't know exactly how reliable each component in an IP network has to be to achieve a given level of reliability (esp. in the face of multiple possible transit paths).

If we knew this information we could more rationally budget our resources to build our networks. We could also potentially design IP networks that are *more* reliable than the telephone network, and run telephony and other more demanded services over them.

But we haven't asked these kinds of questions... so how can we be confident that putting end2end solutions in the network is the right solution???

Mumm: On the Optical Cross Connect front: The usage that may become dominant, is more protocol agnostic. The sales of point to point lambda windows makes for a pretty cool application. This looks more to be a "carrier-to-carrier" application. Rather than leasing dark fiber, you can just buy a lambda to complete your ring...etc. We might finally have a good use for all this glass going in the ground.

Tony Li: A key observation here is that the point of an optical cross connect is to provide a real circuit, not a virtual one.

An optical cross connect, functioning along with IP routing and an intelligent traffic management system can be used to dynamically place bandwidth where it is needed, when it is needed. The optical plane provides an active provisioning fabric, allowing the network to be more efficient. And a more efficient network makes for a more profitable ISP.

I don't see optical cross connects as an opposition to IP technology. Rather, it provides one of the key means of automating the network that is sorely needed. A much better question to ask is: can IP routing possibly survive its projected growth curve without the enabling technology that a flexible optical fabic provides? Yours in dissent, Tony

p.s. Just in case there's any confusion out there, I'm still the world's biggest proponent of IP routing. I just don't assume that we know everything about networking already. I hope that it can be made better.

Bora Akyol: I agree that a more dynamic optical infrastructure allows an IP network to be established faster and better (in terms of flexibility), but I disagree with the point of view that expects routers to dynamically establish, modify and tear-down circuits to other routers on demand. First of all, the current (IGP) routing protocols don't have a clue on who they want to talk to, they talk to whoever is out there and answers their HELLOs. Secondly, we tried this before (ATM) and it did not work.

The current ODSI work has the concept of dynamic provisioning completely upside down, IMHO

Tony Li: The key words in this sentence are "on demand". I believe that we've demonstrated that traffic engineering is a viable and beneficial capability in large scale IP networks. I would agree that anyone attempting to perform traffic engineering with a very small time constant would be pushing the technology past what is beneficial today.

Bora Akyol: 1) I believe that all of the problems that are claimed to be solved by TE can also be solved by a well-designed network architecture and a good routing protocol. Unfortunately for the Internet, the development and research on IP routing protocols that are load-sensitive has come to a halt. I know that there are people working on these including some of us at Pluris, but in general, there is strong pushback on anyone that even |suggests that a load-sensitive routing protocol is a better solution than TE.

Tony Li: You should understand that that pushback is based on a technical history. There have been many experiments with load sensitive link state routing protocols. All of them, from the original days of the ARPAnet, have resulted in instability, with oscillations in the traffic matrix and high CPU loading as all of the routers SPF frequently to keep up.

Personally, I believe that there is a solution buried somewhere in control theory, but the industry hasn't hit on the right human that knows enough about control theory AND routing protocols to make this work. This

has been a pet peeve of mine since about 1987, and yes, everytime someone says that the answer is 'more thrust', we have an educational discussion.

Bora Akyol: If there was a desire to work on such a routing protocol in the community, we would definitely like to help. In the meantime, I will keep on working on such a protocol with a small group of people here.

2) In terms of a management perspective, I think that it is clearly advantageous to manage a single network with no overlay topology. ATM was not even close to this and MPLS although closer still does not meet the goal of the unified network. I am still trying to figure out what exactly is wrong with a combination of fast/dense/scalable routers and optical equipment without an overlay architecture. As an aside, I don't think managing on the order of 5000-10000 LSPs in a core backbone is easy at all.

Toni Li: I don't think anyone is suggesting that managing 5000 of anything is easy. First, you don't need 5000 LSPs to perform traffic engineering. Only enough to redirect traffic away from hot spots. Second, this needs to be automated. This is a small subset of the fact that all of our network management needs automation. Otherwise, we can't possibly hope to get the operator expertise to continue to scale the network.

ICANN FOOTNOTES

Continued from page 26

the move could also be seen as an attempt to pre-empt European efforts to create their own new registry under the .EU or .EUR TLD. A new, truly European TLD would attract much more business registrations from Europe than NSI's proposed .shop, and cut into the dominance of dot com. Shared ownership of the new registry by ICANN-accredited registrars – 60 percent of which are American – would further dilute European market share. The small number of new TLDs under the NSI proposal also shuts out other regions, such as the growing Asia-Pacific region.

Conclusions

If ICANN implements NSI's suggestions, NSI's dominance of the domain name market would be prolonged for another year or more. Management of the domain name space will take on all the features of an international cartel. The NSI proposal offers vested interests privileged access to the new name space while shutting out consumers, non-commercial organizations, and independent entrepreneurs.

Members of the public can submit comments to ICANN on the new gTLD issue at this web site: http://www.icann.org/dnso/new-gtlds-01apr00.htm [Editor: this paper is to be found at: http://dcc.syr.edu/report.htm].

ICANN Footnote 2: ICANN Membership Portrayed to Europeans by Bertelsman as Important Objective

Editor: the Europeans are discovering ICANN. As the material below shows those behind the Bertelsman Foundation Democratic Internet project are far more direct in their assertions about the power that they expect ICANN to have than ICANN staff have ever been. Where they seem to be exceptionally naïve is in taking at face value the assertion that ICANN At Large membership elections can ever provide German or any other internet users with protection against arbitrary and capricious ICANN actions.

On May 8th 2000 European Marc Holitscher <holitscher@ir.gess.ethz.ch> commented on BWG; "As usual, this [Editor: the origin of support for German citizens to join Icann's at large membership group] is not very transparent. The initiative is supported by several other media conglomerates from Germany, France and Austria, the German national TV as well as a representative from the German government. . . . . Der Spiegel" - a very influential weekly news magazine from Germany has started an initiative called "I can! eLection 2000" that aims at increasing German participation in the @large project. The URL is http://www.spiegel.de/netzwelt/icann/ By the way, the Bertelsmann Foundation has just launched its own effort aimed at informing (instructing?) people in Europe about the @large elections. An English version of their website can be found here: http://www.democratic-internet.de/pages/english/home.html"

On the Bertelsman Foundation web site we find the following slippery home page text.

The Internet is changing everyday life: the way we communicate, the way we work, and the way we live. It enables worldwide, interactive, instant communication. However, certain questions still remain.

For example, how is it governed? Who lays down the standards according to which the network of networks is regulated? Who determines how communication on the network takes place?

"One thing is certain: the Internet is largely beyond the reach of national attempts to regulate it. Anew culture of responsibility is developing. Self-regulation of the Net - Internet Governance -appears to represent a promising approach.

It requires individual users to assume anew level of responsibility.

Within the scope of its media-policy projects, the Bertelsmann Foundation is accompanying such self-regulatory processes. The project "Democratic Internet" provides information about the formation of the Internet Corporation for Assigned Names and Numbers (ICANN). ICANN is going to manage key areas of the global network: for example, it will determine how a user obtains an address on the Internet and, therefore, how he is able to shape his identity in the information age. The role of the individual Internet user may be crucial here. The first worldwide online elections to determine the composition of the ICANN Board are due to take place this year.

[Editor: this last sentence is a huge mis-representation of fact. Elections will be held for 5 board seats. But the nominees for those seats as ICANN announced on May 9 will be hand picked by a nominations committee composed of ICANN insiders.]

Then under press announcement we find:

FOR A DEMOCRATIC INTERNET The Bertelsmann Foundation promotes European participation in the ICANN elections.

This year German Internet users will be able to send representatives to an international body which shall decide on the worldwide assignment of all Internet domains and names. ICANN was founded by the US government in 1998. The intention is for it to now become a global organization. ICANN is responsible for the names, numbers and protocols used on the Internet as well as for the root server system which enables the "network of networks" to function.

As Mark Wössner, Chairman of the Board of the Bertelsmann Foundation, emphasized in Gütersloh today: (February 16, 2000) "Whoever has control over the technical roots - the root servers - of the Internet can control how communication, work, reading and buying takes place and, ultimately, how the Internet is governed". The Bertelsmann Foundation's involvement here is explained by a wish to participate in the development of the electoral system and promote a high level of European participation. In the USA, the foundation of ICANN and the worldwide elections to its Board are the subject of broad-based discussions. By contrast, in Germany and the rest of Europe, hardly any attention has been paid to the foundation of ICANN and its fundamental importance for determining the future course of the Internet.

The Bertelsmann Foundation's project

"Democratic Internet" has set itself the goal of cancelling out this deficit. Inter alia, the project intends to provide information about ICANN in an attempt to encourage a broader section of the public and those involved in the political sphere to actively accompany the process. The project shall also contribute towards ensuring that ICANN's organizational structure, its method of working and the course of the elections conform to democratic principles - transparency, representation, equality. The project aims to support the formation of a European ICANN electorate and the nomination of European candidates

Editor: Under project goal we find the following information.

Every Internet user has the opportunity to elect his or her representatives onto ICANN's Board. ICANN will play a leading role in shaping the future of the Internet. The project's goal is to provide information on ICANN's responsibilities and thus help ensure that, in terms of its organizational structure and the way it operates, ICANN upholds democratic principles - transparency, representation, and equality. The Bertelsmann Foundation is accompanying and analyzing the formation of a European ICANN electorate. There is a direct opportunity here for each and every Internet user to decide on his or her own future in cyberspace.

ICANN Footnote 3: ICANN Board Violates Bylaws in Selection of Committee Members And Continues to Work in Closed, Secret Sessions

An essay by attorney **Bret Fausett**, on May 10. (Used with permission and found at http://www.lextext.com/21days.html).

As the heart of its Bylaws, the Internet Corporation for Assigned Names and Numbers pledges to "operate to the maximum extent feasible in an open and transparent manner and consistent with procedures designed to ensure fairness." In fact, so important is this principle to the organization's mission that an entire section of the Bylaws is devoted to "Transparency."

One of those "Transparency" provisions, titled "Access to Information" provides:

"All minutes of meetings of the Board, the At Large Council, Supporting Organizations (and any councils thereof) and Committees shall be approved promptly by the originating body. No later than twenty-one (21) days after each meeting, draft minutes shall be made publicly available on a publicly-accessible Internet World Wide Web site maintained by the Corporation (the "Web Site")...."

ICANN Bylaws, Article III, Section 2. The timely posting of minutes that this provision requires was especially important in the early days of ICANN, when all of the Board's meetings were closed to the public. Even though ICANN has now opened its quarterly meetings to public observation, its monthly teleconferences are still closed. For these closed meetings, the importance of the 21 day disclosure is magnified, as it represents the only insight the outside world has into the operations and decisions of this important body.

And important things happen in these closed meetings. Witness the April 6, 2000 Board teleconference. On April 6th, in a closed teleconference, the Board members first discussed the need to fill slots on the newly created Nominating and Elections Committees. Later this year, 5 new ICANN Board members will be elected, and these two committees are charged with the important task of finding candidates for these positions, placing them on the ballot, and overseeing the election. The Nominating Committee is especially important, as it will determine who runs for these 5 new seats, representing over 25% of the Board.

The minutes of the April 6, 2000 meeting revealed that "Ms. Wilson [one of the current ICANN Board members] will gather suggestions for the non-Director members and present them in advance of the next Board teleconference [scheduled for May 4, 2000], so that the committee charter and membership may be formalized at that time."

Consistent with those minutes, during the May 4, 2000 meeting, the Board selected the members of the Nominating Committee and by May 9, 2000, it had issued a press release announcing the appointment of seven individuals to serve.

If ICANN had followed its bylaws, the minutes from the April 6, 2000 telephone conference would have been published within 21 days, or no later than April 27th. On April 27th, we would have learned that Linda Wilson was "gathering suggestions" for seats on the important Nominating Committee and that the final selection of the committee members would take place at the next meeting on May 4th.

This would have given members of the Internet community an entire week to send any suggestions to the Board. But the Bylaws were not followed.

In contravention of its own Bylaws, ICANN published the minutes of the April 6, 2000 Board teleconference on May 9, 2000 — the same day that it issued its press release announcing the appointment of the members of the Nominating and Elections Committee.

Whether by oversight or intent, ICANN and its staff deprived the Internet community from having any insight into its operations or input into its decision-making.

The magnitude of the error or misjudgment made by ICANN and its staff was compounded by the fact that the selection of these committee members was made in a different fashion than the selection of members for similar ICANN committees. Most of the other ICANN committees staffed with non-Board members — and all of the committees relating to ICANN's At Large membership — were staffed only after ICANN issued a public call for participation. This happened with the Membership Advisory Committee, the Independent Review Advisory Committee, and the Membership Implementation Task Force. It did not happen with the Nominating Committee and the Elections Committee, an important detail running contrary to reasonable expectations based on past history.

— Bret

On May 10 Harold Feld, another attorney, commented:

On reading Brett's essay, I'm inclined to think there may well be a cause of action under California law to enforce the bylaws. All at-large members are potentially harmed, thus creating a pool of people eligible to bring suit.

ICANN Footnote 4: Our Reconstruction of the Events of July 1999

On June 22, 1999 Jim Rutt published his now infamous Rutt Report #1 where he laid out a strong public interest provision for the future of NSI and the Internet. Unfortunately, within months, NSI had (1) been sold to an IBM influenced Verisign, having (2) first sold out by agreeing in return for its own survival to pay the bills of irts arch enemy, and had (3) a former nemesis appointed to its highest policy position. We ask in this essay "What happened?"

From the point of view of the credibility of the new NSI CEO, the most damning thing that happened was that, having told the internet community that he respected it, understood it, believed in the need to maintain its freedom and culture and then having promised to take a continuing dialogue with the community, he fell silent. He issued no

more reports and even failed to do such a rudimentary thing as appoint someone to the position of customer ombudsman at NSI as he had promised he would. We had conversations with him about his goals for NSI in June of 1999. They were entirely compatible with the intentions he outlined in his Rutt Report. Those to whom we have talked feel that he was sincere about the pact that he made with the Internet on June 22. Yet during he next few weeks he abrogated his solemn pledges and, under his "ostensible" leadership, NSI has gone one to become the very opposite of the company that he pledged to make. Rutt now "leads" an organization whose customers basically have no rights to the Internet addresses for which they pay.

NSI's current idea of market leadership is to do anything to churn the turnover of its product. On May 10th it stated in a press release: "Beginning today, customers can fill out a simple form on Network Solutions site, and NSI will list domain names in .com, .net and .org on its website for customers who are interested in marketing their domain names. Network Solutions will provide the service for free to its customers for an introductory period." The press release failed to warn NSI's own customers that by placing domain names up for resale they could be nailed under the ICANN UDRP policy by an entity wanting to have the domain name awarded to it for bad faith speculation in someone else's name. As an observer commented: "Gee. Register your name for sale with NSI and provide reverse hijackers with their evidence of bad faith in documented form! What an idea!"

We ask here how Rutt could have managed to make such a mockery of the personal integrity that he presented to the Internet in Rutt Report #1? What happened to cause him to abandon the very goals in which he swore that he believed? We now have seen some indication of what happened during the first days of July of 1999. NSI was moving forward on at least three different fronts. Work on Rutt Report #2 was underway. Drafts by multiple authors were under review. We have a copy of one such draft. Calm cool short and to the point. Independent. Not especially threatened or threatening.

Rutt was preparing for his Congressional testimony on July 22. It is not clear who was advising him. We reached him by phone in his office at 6pm on July 23 the day after the debacle and asked him whether he agreed with the comment that he had his head handed to him on a platter up on the hill. He admitted that that was an accurate characterization and blamed it on "handlers" brought in to help him prep. He also admitted that in a departure from customary procedure top NSI attorney Phil Sbarbaro did not brief him. The consensus that we have gathered from conversations with many sources is that he did not take the occasion seriously and perhaps thought he could charm the congressmen by appearing as "likeable" country bumpkin. His ex post facto "spin" claims that he was coming on wild and unpredictable so that the congress critters knew they better settle on his terms. The reality of his not testifying at the intellectual property hearing the following week shows this spin to be wishful thinking on his part.

Parallel with Rutt Report #2 NSI was building and deploying alternate root servers. Chris Clough was heard to brag to NTIA's General Counsel Andy Pincus that NSI would deploy alternative servers. Rutt himself has bragged that they had done so and moved them outside the NSI firewall telling DoC. what it had done. David Holtzman has stated the same thing adding that DoC sternly warned that, if they were used, the official servers would be effectively nationalized by the administration. The administration would declare them a "strategic telecommunications resource" and would presumably move the legal and operational responsibility for operating them from NSI to the Department of Commerce. The gauntlet was very likely thrown down that the press would be told that NSI's self serving actions would be presented by the US government as giving it no choice but to impose regulation on DNS. Furthermore that such action would mean the failure of the ICANN self-regulation strategy. Moreover a US move to regulate DNS would be seen as likely to lead to general regulation of the Internet and would slow down the growth of electronic commerce on which NSI's future prosperity depended.

In trying to reconstruct what happened it would be extremely useful to understand three things. 1) Why Rutt #2 never got out the door. 2) When the administration's threat to NSI on the root server issue was first made. 3) Precisely what was threatened. It is very likely that in the week before the hearing, the game of dare and double dare on the root server had begun. NSI's strategy may have been to force the administration hand pending a successful hearing on July 22. Based on what we can ascertain, the Rutt #2 drafts were not confrontational and alternative root server oriented. Putting one out in such a tense situation would have sent the wrong signals. But when Rutt fell on his face in front of Congress, NSI's ability to face down the administration threat evaporated. The only road open was a settlement with the enemy. That path was embraced at the July 31, 1999 meeting attended by Farber, Kahn, Bradner and Jim Rutt's close friend Brian Reid.

We cannot prove that events happened precisely as we have speculated they did. However for every assertion made above, we have done so either on the basis of our own direct knowledge or on conversations with trusted sources who claim fist hand knowledge of their own. This essay has been our best attempt based on many conversations with primary sources and reviews of several leaked documents to piece together an understanding of otherwise inexplicable events. Any further assistance in helping to fill in the gaps will be appreciated.

Footnote 5: Network Solutions Security Lapses Continue

Editor's Note: With ICANN assigned the role of ensuring the stability of the DNS Network Solutions is freed to take a cavalier attitude towards security. IF it gets caught, it merely points the finger at ICANN. On May 9th a user was ale to grab a CGI script that allowed any file on the Network Solutions secure systems to be read. He posted the URL below to NANOG. The URL retrieves a file of the log n ID's of every NSOL employee with access to its secured systems. The NANOG post, the file contents and an explanation of what had happened all follow.

Date: Tue, 9 May 2000 13:53:54 -0700 (PDT)
From: Exiled Dave <exiled_dave@yahoo.com>
Subject: NETWORK SOLUTIONS SECURITY TRUTHS?
To: nanog@nanog.org

http://www.networksolutions.com/cgi-bin/makechanges/easysteps/easystepsSIRN=mynumSIRN=mynum&FILE=../../../../../../etc/passwd

Check out /var/named/internic.net as well.

[Editor: We grabbed the contents of this UR and reposted then to BWG]

root:x:0:1:Super-User:/:/bin/csh daemon:x:1:1::/:/bin:x:2:2::/usr/bin: sys:x:3:3::/: adm:x:4:4:Admin:/var/adm: lp:x:71:8:Line Printer Admin:/usr/spool/lp: smtp:x:0:0:Mail Daemon User:/: uucp:x:5:5:uucp Admin:/usr/lib/uucp: nuucp:x:9:9:uucp Admin:/var/spool/uucppublic:/usr/lib/uucp/uucico listen:x:37:4:Network Admin:/usr/net/nls: nobody:x:60001:60001:Nobody:/home/nobody:/bin/csh noaccess:x:60002:60002:No Access User:/: nobody4:x:65534:65534:SunOS 4.x Nobody:/: rholgado:x:1063:14:Ruel Holgado:/home/devel/rholgado:/bin/csh abolivar:x:736:14:Alejandro Bolivar:/home/devel/abolivar:/bin/csh shanes:x:554:14:Shane Smith:/home/devel/shanes:/bin/csh thuann:x:531:14:Thuan Nguyen:/home/devel/thuann:/bin/csh jcao:x:1004:10:Jin Cao:/home/devel/jcao:/bin/csh rturner:x:751:14:Rodney B. Turner:/home/devel/rturner:/bin/csh vlee:x:734:10:Vinny Lee:/home/devel/vlee:/bin/csh bjoseph:x:904:10:Bernard Joseph:/home/devel/bjoseph:/bin/csh www:x:9999:10:www:/app/http_sicon:/bin/csh dummy:x:10000:60001:OLPS:/home/devel/dummy:/bin/csh regdb:x:60:10:Registration Database:/home/reg/regdb:/bin/csh dumpuser:x:222:1:Amanda Dump

Continued from page 36

Executive Summary

Hughes on Wireless pp. 1 - 15

We interview Dave Hughes, principal investigator of two NSF wireless projects and owner of Old Colorado Communications on the state of the art of TCP/IP radios for wireless ISPs and for scientific environmental field research (his second *major* NSF project). He points out how a fall in prices and increase in capability has pushed the price of 10 Mbps radios to well under $1000 each.

He views Cisco's purchase of Aironet as a major move forward for the wireless spread spectrum industry however the impact of this will ultimately depend on how Cisco integrates Aironet products into its business line and whether it comes to see itself as a provider of connectivity solutions. Cisco has announced its own 45 megabit per second LMDS radio and is also bringing out a line of UNII band radios limited to five miles in transmission range.

These new radios can be remotely logged into and configured - something that greatly increases their utility for ISPs. ISPs meanwhile are going wireless. Breezecom claims ISP 500 customers in North America. Some 100 megabit per second radios are beginning to appear. One such is made by Proxim.

While line of sight problems are critical for these radios, Cisco is claiming to have overcome some of the drop off of communications caused by Fresnel Zone problems. Subtle physical differences found at each site can mar transmission capability. ISPs must have staff skilled in installation. The key business model is focused on connecting small business and will be increasingly focused on delivery of broadband services to residential customers who either don't want or can't get adequate DSL or cable connection.

Hughes discusses in detail the way that the E-rate increases by an order of magnitude the cost of connecting public schools to the Internet by prohibiting the schools from buying wireless equipment and requiring them instead to rent leased lines from the local exchange carrier year after year. Under the Texas subsidy, the monthly cost of the Internet connection is multiplied by 12 and to that figure is added the cost of hardware (DSU/CSU for example) needed for the phone connection. The total sum becomes the amount of subsidy for which the district is eligible. The district is then free to spend the money on the telco connection or on radios and a wireless connection. Wireless normally wins because the district after costing out the alternatives, normally finds that the cost of radios and plug in via radio to an ISP leaves them several thousand dollars left over.

At the 2.4 giga hertz range most manufacturers make radios that operate at one tenth of a watt or 100 miliwatts rather than the allowed power of one watt. They do this because [such radios can be sold in Europe, and] it saves considerable money in the cost of manufacturing. Customer don't seem to mind because if their cheaper 100 miliwatt radio won't connect they can buy an inline amplifier for $750 and increase the power to a full watt. When this has been done and under ideal line of sight conditions the radios have successfully work over distances of up to seventy miles.

Given the lack of incentive for wireline telco's to bring broadband into rural America, the FCC is has issued a notice of inquiry on the subject of a Software Defined Radio (SDR). One where smart software controls the radio - its power, its frequency spread, and other technical characteristics. Major spectrum possibilities could be achieved simply by allowing the design and use of radios that could tune themselves in accordance with the operational reality of their surroundings.

Hughes points out that one way a user who lives close to an ISP pop can help to spread the benefits of wireless technology is to ascertain whether the ISP operator with allow him to plug a radio into the ISPs pop connection. If so after a [site] sight survey to determine that radios needed are available at reasonable price, the user can install the radios and plug one into the ISP's Ethernet.

To ensure that they can inter operate with each other, radios are being built to the 802.11 standard. As shown by Apple Computer in its Macintosh Airport Base Station and Airport card architecture, it is possible for someone to spend upwards of a thousand dollars to connect to an ISP with a point to multi-point multi megabit radio and then by placing an omni directional antenna on his roof to connect and relay as many as twenty neighbors using $100 PCMCIA radios plugged into their lap tops.

Of special benefit to ISPs selling wireless connectivity is the ability in software to throttle down to agreed upon rates, the speed of the connections that radios hey supply to their customers give. Also of significance is a new Ethernet PPP protocol that means they need not assign each customer their own IP number.

"Watch out UUNET, the frogs and the shrimp are coming - using your bandwidth" - Dave Hughes

Having laid out this general background, Hughes goes on to discuss aspects of his current NSF funded research on wireless and satellite connectivity for environmental research.

Hughes emphasizes Tachyon which has just come on line in the spring of 2000 with Concentric as its Internet providing partner as the first company to provide reasonably priced bi-directional satellite linkage to the internet. Tachyon provides a ground station that talks to the satellite for only $5,000.. It sells bi-directional, true TCP/IP, at 2 megabits down and 256k up, for $2000 a month, or 300kbps down and 64k up, for $795.

Hughes describes the National Environmental Observatory Network as part of an expanding need for environmental data collection — one that is so broad in its proposed scope that it looks as though only wireless data monitors may do an cost effective job of data gathering.

He talks about several projects in which he is working with environmental scientists whom he is surprised to find are generally unaware of the what improvement in wireless data gathering technology over the past five years make it possible to do.

The kinds of data gathering involved are quite diverse. For example the transmission from sensors the chemical composition of lake water in timed coordination with the overhead passes of a satellite. A second is the collection of light readings from a network sensors on a forest floor and the transmission of that data from each individual sensor a short distance to a data collector. The collector, in turn after perhaps encapsulating it in a tiny Linux data base, transmits it back to the research station and the Internet. A third is the capture of sound - in this case the call of the coqui frog from the rain forest of Puerto Rico. A fourth is the visual observation of shrimp transmitted in high bandwidth in real time. A possible fifth would include the use of partially buried sensors to grab, chemically analyze via tiny Linux and transmit the composition of the gasses of a prairie fire as the fire ignites the material around the sensor and passes over head. Other phenomena measured may be as diverse a earth quakes and hurricanes. From the point of view of Internet infrastructure this means yet another huge increase in bandwidth that will be generated and dumped on to backbones. Watch out UUNET, the frogs and the shrimp are coming - using your bandwidth.

Rosenberg on IP Telephony, pp. 16- 22

We interview Johnathon Rosenberg Chief

continued on next page

Continued from page 34

Scientist of dynamicSoft. Rosenburg updates us on the outcome of the Megaco protocol which is designed to facilitate the communication of IP telephony gateways with SS7 switches in the Public Switched Telephone Network. The ITU successfully cooperated in the development of Megaco and in January 2000 held a meeting with in Geneva with a small handful of the top IETF leadership. [**Editor's Note** and not part of Rosenberg interview: Reports from attendees at the meeting indicate that for he first time ITU leadership was overtly eager to cooperate with he IETF in completing the protocols necessary to achieve converge of the voice and data networks. It was also interesting to note in view of the evolving relationship between the IETF and ICANN that Karen Rose, the number two person to Becky Burr also attended.]

Among the other protocols that Rosenberg describes are ENUM, PINT, SPIRIT and SCTP. Here the purpose begins to be to have a users computer link with the PSTN and initiate events there that formerly could have been done only through the network's intelligent switching system. While Quality of Service issues are still unresolved and are critical to those who would merely move voice telephony to the Internet, Rosenburg and dynamicsoft advocate a blending of internet and telephony capabilities. Rosenberg points out that with every new medium added to the mix the number of applications enabled grows exponentially. Thus voice telephony over the internet is nothing more than the transfer of a standard century year old service. Internet, video, and voice raises a host of new possibilities. Add the web to this and the opportunity for flexible and powerful productivity enhancing tools seems vast.

Using the SIP protocol dynamicSoft specializes in the provision of client server based tool kits that can be tailored to the needs of individual companies. Under these conditions it becomes possible to think in terms of where the arrival of specified kinds of email could trigger application with the telephone network on behalf of the user. Rosenberg sees a converged future where the only telephony that is worth having is web-enabled.

Obstacles to Bandwitdh Commoditization, pp. 23- 26

In Part 3 of an on going series we interview Noel Chiappa developer of the first multi-protocol router. The interview focused on issues involved in the development of a commodity exchange for bandwidth. While we speculate on the changes in the power structure of the Internet industry that this is likely to bring on, Noel points out that the thing most likely to slow the commoditization of bandwidth is "that Internet routing, isn't ready for it yet. It's all spit and bailing wire. . . . . Commoditization implies a tremendous amount of flexibility, a high rate of change in topology in the way things are connected together. It also requires a very robust infrastructure. That's one thing we certainly do not have. A lot of this stuff runs because there's a lot of smart people tweaking it all the time. And I'm not sure we have enough smart people." For example when one buys an OC3 from San Francisco to Atlanta to start at Sunday midnight and run for 24 hours one is going to want to be certain that the interfaces with the rest of the Internet adjust smoothly. To do this: "You want the routing to adjust very quickly so that you can use it and when it goes away, you want the routing to adjust back. And that's my question: can the routing really do it?" "If you start dorking around with the topology of the network, are the protocols and everything else in it robust enough to deal with that kind of sort of brownian motion in the connectivity?"

Engineering issues; IPv6, NAT, IP Telephony Conundrum, Optical Cross Connect pp. 27 - 31

On the IETF, NANOG, and Inet-Access lists the IPv6 versus IPv4 and NATed-end-points religious wars continued. It has become fairly clear that the inertia favoring the continued use of IPv4 is vast. A change over would be hugely expensive and no one seems to know quite how to incentivise network operators both on the back bones and end points to do it. Suggestions as extreme as have the federal government mandate a switch were heard. Apart from a desire a desire to salvage protocols like IPsec that were written with end-to-end network transparency in mind, there was the growing realization that many of perceived benefits of IP telephony could not be achieved IP address translation devices (NAT) boxes stood in the way. Richard Shockey put it bluntly "The deconstruction of the PSTN will be impossible without the introduction of IPv6 and the elimination of NAT's in private networks." We present some highlights of the discussions.

Discontent with the policies of the routing registries is growing. Sean Doran presented a useful critique of the current policy. An excerpt — This is a system which enforces a "one-seller" (the IANA), "one buyer" (one may return addresses to IANA only) model, which flies in the face of free markets, and perversely imposes costs upon consumers.

Although I am happy that there are people trying to conserve IPv4 addresses and also encourage sensible routing announcements by providing not less than a sizeable aggregatable range to qualified buyers, the qualification process is tricky and gets trickier as one's business grows.

Finally on NANOG Tony Li had a few things to say about optical cross connects. "An optical cross connect, functioning along with IP routing and an intelligent traffic management system can be used to dynamically place bandwidth where it is needed, when it is needed. The optical plane provides an active provisioning fabric, allowing the network to be more efficient. And a more efficient network makes for a more profitable ISP."

ICANN- NSI Footnotes, pp. 15, 22, 31 - 33, 36

We continue to document some of the more egregious actions of ICANN and Network Solutions. As the first footnote we republish Milton Mueller's April 25th showing the ridiculous self-serving nature of Roger Cochetti's comments in favoring the introduction of two new to level domains. Mueller makes fours points and then goes on to under score each one in crisp detail.

"1. It would require the new (shop) registry to offer exactly the same terms and prices as the NSI com/net/org registry 2. It drastically limits the number of competing registries, for no good reason. 3. Its ownership arrangements would institutionalize cartel-like controls on the name space. 4. It would put NSI in charge of the back-office services of one the .banc registry, further reinforcing NSI's dominance of the domain name registry market."

The second footnote is about the Europeans who are beginning to discover ICANN. Some large corporate content forces there have mounted campaigns designed with the deluded hype that European participation in the ICANN at large membership process can give Europe a role in ICANN's regulation of the Internet. More information may be found at http://www.democratic-internet.de/pages/english/home.htm

For those who have watched how ICANN operates in ignoring the wishes of all its working groups the following passage seems just a tad misleading "One thing is certain: the Internet is largely beyond the reach of national attempts to regulate it. A new culture of responsibility is developing. Self-regulation of the Net -Internet Governance -appears to represent a promising approach."

The third footnote is a May 10th essay by Brett Fausett titled: "ICANN Board Violates Bylaws in Selection of Committee Members And Continues to Work in Closed, Secret Sessions". Brett explains how the ICANN Board in a meeting on April 6th, that in violation of ICANN's bylaws was not disclosed until May 9th, maneuvered the announcement of a technical elections committee and a nominating committee in such a way as to achieve a fait accompli and deprive the Internet community of all opportunity for input into committee membership.

We have authored the 4th footnote which presents some new information on the events of July 1999 at Network Solutions. We offer a hypothesis of why Jim Rutt broke his pledge made to the Internet in Rutt Report #1 on June 22, 1999.

The 5th Footnote documents NSI's security lapse May 9 when a user found a CGI script that allowed anyone to read any file on NSI's secure systems. the user turn the script into a url that pulled down the names of NSI's secure servers and the logins of their authorized users and log ins. We reprint the contents of what the script retrieved

Continued from page 33

User:/home/dumpuser:/bin/csh akbart:x:591:10:Akbar Tokhi:/home/devel/akbart:/bin/csh pezoua:x:853:301:Paul Ezoua:/home/devel/pezoua:/bin/csh jasons:x:542:14:Jason Stone:/home/devel/jasons:/bin/csh chriso:x:526:14:Christopher Owen:/home/devel/chriso:/bin/csh yosephs:x:2298:14:Yoseph Shiferaw:/home/devel/yosephs:/bin/csh pvirador:x:2253:14:Peter Virador:/home/devel/pvirador:/bin/csh ccurtis:x:772:14:Chris Curtis:/home/devel/ccurtis:/bin/csh jwu:x:878:14:John Wu:/home/devel/jwu:/bin/csh dbrimber:x:2292:301:Donald Brimberg:/home/devel/dbrimber:/bin/csh tlangtry:x:875:301:Travis Langtry:/home/devel/tlangtry:/bin/csh nicstats:x:1100:10:NIC Statistics:/home/devel/nicstats:/bin/csh kmarshal:x:2333:10:Kristen Marshall:/home/devel/kmarshal:/bin/csh afasano:x:2373:10:Anthony Fasano:/home/devel/afasano:/bin/csh hchu:x:3016:10:Hong Chu:/home/devel/hchu:/bin/csh tjohnso:x:3015:10:Thomas Johnson:/home/devel/tjohnso:/bin/csh randrews:x:1076:10:Robert Andrews:/home/randrews:/bin/csh schauhan:x:877:10:Sanjeev Chauhan: /home/schauhan:/bin/csh prepay:x:50000:50000:Prepay batch account:/home/prepay:/bin/csh jbrooks:x:300:10:Jennifer Brooks:/home/devel/jbrooks:/bin/csh cporter:x:307:10:Chris Porter:/home/devel/cporter:/bin/csh pfaber:x:308:10:Pat Faber:/home/devel/pfaber:/bin/csh acarver:x:309:10:Art Carver:/home/devel/acarver:/bin/csh kshepard:x:335:10:Kent Shepard:/home/devel/kshepard:/bin/csh jzeits:x:311:10:John Zeits:/home/devel/jzeits:/bin/csh sduvall:x:330:10:Sean Duvall:/home/devel/sduvall:/bin/csh saictiv:x:99999:10:Saic Tivoli:/home/devel/saictiv:/bin/csh idnames:x:60000:10:idNames:/opt/idNames/home:/bin/csh jcivitar:x:361:10:Jim Civitarese:/home/devel/jcivitar:/bin/csh mlee:x:2282:10:Michelle Lee:/home/devel/mlee:/bin/csh cgreen:x:993:10:Craig Green:/home/devel/cgreen:/bin/csh jwest:x:996:10:John West:/home/devel/jwest:/bin/csh kangell:x:371:10:Karena Angel:/home/devel/kangell:/bin/csh bnesbit:x:379:10:Brian Nesbit:/home/devel/bnesbit:/bin/csh matthewh:x:908:10:Matthew Ho:/home/devel/matthewh:/bin/csh dberry:x:409:10:Dan Berry:/home/devel/dberry:/bin/csh skenneth:x:341:10:Swanson, Kenneth:/home/devel/skenneth:/bin/csh louied:x:376:14:louied:/home/devel/louied:/bin/csh jmcinnes:x:462:10:John Mcinnes:/home/devel/jmcinnes:/bin/csh

List members explained: That is a CGI script that existed so you could make changes to your domain there. The unfortunate side effect was, this script didn't constrain what files you could read and if you mucked around with the URL you could read ANY file off that UNIX server. ANY file. Given a list of all the userids on the system and a list of trusted hosts, it'd be a tad bit easier to hijack NSI and all that it surveys. The passwords were shadowed, so they weren't present in the file; if they had been it would have been truly trivial to walk right into their network with superuser status. Look for lots of slowdowns over the next few days while NSI redoes every account on its network.

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