



# Wireless Reaches Internet Critical Mass

## Internet Use Jumps to Mobile Platforms as Spread of Digital Infrastructure Enhances Wireless Capabilities

### We Survey Issues and Players in Internet Wireless Services

**Editor's Note:** Ira Brodsky is President of Datacomm Research Co., a St. Louis area technology and market forecasting firm. Brodsky has written two books on emerging wireless communications. His firm also published the report *Portals to Profit*, reviewed by *The COOK Report On Internet* earlier this year. We interviewed Ira Brodsky on December 23.

**COOK Report:** What is the wireless "big picture"? Is it that just about all the activities which are being engaged in on the Internet from wired platforms are currently in the process of moving to wireless platforms? Is it also that these wireless platforms represent huge amounts of additional bandwidth consumption because wired platforms are generally not being abandoned for wireless?

**Brodsky:** Yes. That is a reasonable assumption, although a bit of an Internet centric point of view. If I were speaking from a purely wireless centric point of view, I'd describe it a bit differently. In this sense, from an even bigger picture view point, what we are seeing now is the transition of wireless devices from analogue to digital. As a result, for the first time, wireless technologies are able to reap the benefits of Moore's law.

**COOK Report:** In other words while voice was the dominant form of telecommunications, it is not surprising that wireless started out dominated by analog technologies. However, as the bandwidth of digital communications now exceeds that of voice, it is no surprise to see wireless pushing new digital technologies.

**Brodsky:** Correct. What is happening now is pushing the envelop to achieve a collection of new features, performance improvements and cost reductions. It used to be that

this stuff was all analogue and, as such was moving in slow motion. There was just not that much you could do to improve an analogue mobile phone. Now with digital taking over amazing feature sets are being ported to these little phones. For example you have phones now with built in GPS receivers and micro browsers.

**COOK Report:** So how then should a dedicated observer of Internet physical infrastructure analyze developments in the wireless world?

**Brodsky:** You need to understand that something is going on in the wireless realm that parallels what has been going on in the wired world. This something is the explosive growth of bandwidth. It has been accomplished a little bit differently but, due to a move to higher frequencies, to lower power and to more efficient digital technologies it has happened. Five years ago even most radio engineers believed that wireless was doomed to remain low bandwidth. But today the picture is very very different. You have radio systems running at 100 megabits per second with a clear path of future development to speeds as high as 622 megabits per second.

**COOK Report:** So how do they achieve those higher speeds?

**Brodsky:** When you go up higher in the radio spectrum, you have more room to transmit. If you look at the AM radio band, you will find that it is down at the 500 kilohertz to one point five megahertz range of the spectrum. Down there you obviously don't have a megahertz of bandwidth to play with since the whole range of the spectrum application is only about a megahertz. Therefore individual stations transmitting in that band can only have tens of kilohertz. However, if you go up to 30 to 40 gigahertz

Volume VIII, No. 12, February 2000  
ISSN 1071 - 6327

of spectrum, you have the possibility for each user to have hundreds of megahertz of spectrum.

### Broadband Wireless

**COOK Report:** So the speed is a factor of the breadth of bandwidth that one can occupy at the higher end?

**Brodsky:** Yes. However there are also ways that one can make more efficient use of spectrum in lower frequency ranges. For example in the cellular telephone band which is in the 800 megahertz range or in the PCS band which falls in the 1.9 gigahertz spectrum range, you can make more efficient use of bandwidth. You can then allow each user to use more bandwidth.

**COOK Report:** How is this accomplished?

**Brodsky:** In the very high frequencies like 30 to 40 gigahertz, it is generally done by using time division multiple access (TDMA). In the cellular and PCS bands you have both TDMA and CDMA (code division multiple access) which is based on spread spectrum.

### On the Inside:

Wireless Survey	pp 1 - 6, 10
State of Internet 2000	pp. 7 - 10
IETF Architecture	pp. 11 - 24
Farber to FCC	p. 24
Executive Summary	pp. 25-26

**COOK Report:** CDMA would be amenable to increases in its capabilities as a result of continued improvements in chip design whereas TDMA would have fundamental limits in its ability to be improved?

**Brodsky:** Correct. With TDMA you face a fundamental problem. When you move to higher speeds, you have to narrow the time devoted to the transmission of each bit over the airwaves. The biggest problem in any wireless environment is what is called multi path interference. Narrowing the intervals between the transmission of each bit in order to increase the speed of transmission increases your multi path interference problem.

Now with CDMA you are using codes instead of time slots to distinguish one user from another. With CDMA systems in use today in the PCS and cellular systems, each user is given a 1.25 megahertz wide channel. With TDMA users we are restricted to channels in the 30 kilohertz range. With direct sequence CDMA you spread the signal out over the whole range. The data from multiple users can be basically combined into a single transmission. And the code of each user can be used to filter back out the information intended only for that user. If you look at an entire range of CDMA transmission, you see a kind of composite stream that amalgamates into one continuous transmission the contributions of all users of CDMA technology who are using it in any single given area. CDMA literally puts multiple users on the same frequency at the same time.

**COOK Report:** Didn't George Gilder make reference to the these separate codes as languages?

**Brodsky:** I think the analogy that he used was the cocktail party where conversations are going on at the same time and you can by listening for the timbre of familiar voices distinguish one conversation from the others. In CDMA you encode the conversation at one end and then use the same code at the other end to filter it back out.

Normally in a digital world we consider every bit to be precious and it is understood that, if one bit goes wrong, the whole transmission may have to be resent. In CDMA however, you are expanding the data. You are turning each bit into smaller units called "chips". These "chips" are no longer so precious. In fact you work with them on a statistical basis. Therefore, instead of caring desperately about each and everyone, you look at 127 that go by and, as long as the majority of them are received correctly, you get the right user data.

Another nice feature about CDMA is that it

is very elastic in terms of capacity. If you take a CDMA cell with only a handful of users, the geographical coverage of that cell will be at its maximum. If you load the cell with more and more users the coverage of the cell starts to shrink. This feature will permit coverage of rural areas with fewer cells and hence less expense than are necessary to install for a CDMA system in an urban area.

## Wireless Internet Services

**COOK Report:** so having pointed these developments out as making broadband wireless feasible, if one wants to begin looking at wireless internet services, where does one turn? Can one divide up the field into several different topologies? Say fixed services, wireless LANs, mobile services? But then with each of these you have to look at the amount of bandwidth that you are trying to deliver to individual users and whether you are trying to get all users always on connections?

**Brodsky:** Yes. We need also to be looking at one other factor. Namely whether they want any degree of portability. Mobile technology almost always converts easily to fixed but certain types of fixed are not suitable for mobile.

**COOK Report:** So how about giving an inventory of the players in point-to-point?

**Brodsky:** One has to qualify one's answer to your request. Up until a couple of years ago I would have said it is all either point-to-point or point-to-multipoint. But now we have an interesting new player doing what it calls point-to-point-to-point. It is a company called Triton Network Systems in the Orlando, Florida area. They make equipment that operates in what's called the millimeter wave bands. We are talking about frequencies at 28 giga hertz the LMDS or Local Multipoint distribution service. There is a similar band at 38 giga hertz. The focus of this company is on providing basic high speed connections. Therefore they are not focusing on the small office or the home office. What they offer is a very intriguing alternative to fiber. They connect buildings to each other at distances of 2 to 3 miles.

The beauty of this is that if you take an urban area with a lot of buildings where you can cover two miles or more between buildings you can interconnect the whole city. Each building node in your network becomes a relay point by which you may go to the next building.

**COOK Report:** This sounds a bit like part of Next Link's business plan.

**Brodsky:** Next Link does own a great deal of LMDS spectrum in which Triton is lurking. But some competing technologies exist. These are point-to-multi point technologies. Some of which are frequency division and some of which are time division multiplexing. Thus there have been advances in our ability in fixed environments to go to higher speeds using time division multiplexing. When I had mentioned before that TDM was kind of at a dead end, I was referring primarily to its use in mobile technologies. In fixed location use you can squeeze some extra life out of it. In a fixed environment multi path interference won't be as bad because users aren't continually in motion. If you are not moving, you can also tune the system to minimize multi-path interference. You will try for a line of sight signal and you will filter out as many reflections as possible.

**COOK Report:** One of the major marketing points for Triton in cities must be to avoid the expense of digging up streets.

**Brodsky:** When one looks both at the time and the cost of running fiber to an existing building, one will find a technology like Triton quite attractive. The total cost for radios and installation of all other equipment for linking each building to the network is about \$45,000. The bandwidth is 100 megabits per second per link. The equipment involved is a small box of a little more than a cubic foot in size that is bolted to the side of the building. The box is pointed towards the next node in the network. You may want to have 2 of them on a building. One in a position of receiving from the building down stream and the second in the position of becoming the transmitting device. Note that if you connect several buildings you can have redundancy because you have created a ring, which should it be broken at any spot can be traversed in the other direction.

At the millimeter wave length the radio waves behave somewhat like light. You can view the transmitters as being something like searchlights. You actually have to point them at the receivers on the next building. But one other benefit of this aiming is that you can have transmissions criss-crossing each other without interference.

Now another big advance is the aggressive reuse of frequencies. The more you reuse frequencies, the more wireless begins to provide the same benefits of wired. In one sense you can think of wired as being the ultimate wireless system in that it allows you to reuse frequencies as close together as you wish because the frequencies are confined within the wire. In the wired world you can reuse the bandwidth of a single wire from point "a" to point "b" by running a bundle of a hundred wires between "a" and "b". In the

wireless world by using higher frequencies, lower transmit power and more efficient digital technologies, you can reuse frequencies more aggressively and therefore get more capacity and more bandwidth.

**COOK Report:** You have more means of isolating your transmissions and more flexibility in doing so?

**Brodsky:** Exactly.

**COOK Report:** So what would your overview of mobile look like?

**Brodsky:** The big development here in 1999 was simply the huge growth of mobile telephones world wide driven by transition of these phones from analogue to digital technology. As I said earlier, all the benefits of Moore's law are being brought to bear on the capabilities of mobile phones for the first time. The transition from analogue to digital has meant that mobile data capabilities are under going an explosion of development.

Editor's note as we go to press: We note that one of the most powerful and popular mobile digital data devices is Qualcomm's pdQ phone, which has a built-in version of the PalmPilot handheld. In August 1999 Sprint announced Sprint PCS Wireless Web, with services available in late September. One part of the service is wireless Internet access. Users can either access the Internet directly with CDMA phones including the Qualcomm pdQ, or use the phones as wireless modems and connect them to devices such as handhelds and notebook computers.

According to the December 1999 *Gilder Technology Report*, on November 9 Qualcomm demonstrated a new 2.4 megabit per second burst capable High data Rate (HDR) system compatible with all existing CDMA deployments and offering a range of services including internet access. Ironically HDR is a dynamically adaptable form of power controlled TDMA which fits seamlessly into unused 1.25 megahertz CDMA channels. With HDR coming on in 2001, you will be able to plug your notebook into your CDMA cell phone and get faster access than your office T-1 line. Soon the HDR chipset will become your PC's on board wireless modem." This will become part of the additional tsunami of wireless internet access with which internet architecture will be forced to cope. [According to Brodsky: "HDR is CDMA; there may be some time-sharing going on, but it would be misleading to call it "TDMA." It would be more accurate to say HDR runs on separate channels that can be either in the cellular/PCS spectrum or outside that spectrum; saying "unused channels" suggests it borrows channels from the voice system." ]

## Mobile — Web Phones

**COOK Report:** Well from the Internet data point-of-view are the web phones most important?

**Brodsky:** That is indeed the most likely case. There are a lot of different parties involved in making this happen, it looks like Phone.com which invented the basic concept of putting a micro browser in a phone, and then delivering content from the Internet to those micro browser equipped phones. Originally Phone.com was named Libris. Then they changed their name to Unwired Planet and finally to Phone.com. Their founder is a French gentleman named Alain Rossmann. He was the founder of EO Computers. Now they made a tablet computer and then a pen computer to which they added a cellular module. At this point ATT bought the company only to fold it when the company only sold 80,000 of the EO computers during its first year of life.

But Rossmann really learned his lesson. By making a wireless tablet computer he realized that it was too much for end users to swallow. What you really needed to do was not turn the computer into a combination computer and cell phone. Instead he decided to take the basic cell phone platform and add just a small amount of data capability to it and see what he could do with it. He learned immediately that he could do a lot of useful things like white page look-ups and even on line trading. In December 1999 there has been a series of full page ads run by AmeriTrade in the Wall Street Journal showing the Sprint Wireless Web Phones with a stock trading menu.

**COOK Report :** Is a micro browser there to give the user some of the minimal capabilities of its full-fledged cousin?

**Brodsky:** You are talking about an application that allows you to view a limited amount of text. It strips out all the graphics. Furthermore the amount of information that can be shown on the screen is also very small. The micro browser is making it feasible to present the mobile users with these small pocket sized devices with key information that they can use. Phone.com invented the micro browser. It has also lead to the development of a standards forum called WAP which is designed to formulate a standard for micro browsers to adhere to. Although it may cost a lot of money to join the forum, WAP will be an open standard that anyone can use. While there are a number of companies making their own micro browsers, most of their products are WAP compliant.

**COOK Report:** Then, the purpose of WAP is to have a standards framework for micro browsers?

**Brodsky:** I think you can have completely different micro browsers and still be WAP compliant. Nokia has its own. And Microsoft has announced a deal with Ericson to provide a micro browser for them.

**COOK Report:** So what then does WAP compliant mean?

**Brodsky:** WAP is the protocol. It says here is how we are going to send data. How we are going to format it for delivery to a phone. It is the language for the micro browser. Just as HTML is the language for a web site. For example Netscape and Microsoft Internet Explorer are two different browsers that both use HTML.

But WAP is really more than a language for a micro browser. It is a full protocol stack for use in accessing the Internet from small devices with limited input and out put capabilities. These can be PDAs as well as Web phones. WAP includes security features as well as transmission and content issues. It covers a wide range and not everyone is expected to be fully WAP compliant in all areas.

**COOK Report:** When you say content issues to what are you referring?

**Brodsky:** They have come up with something called the wireless markup language. When they created the WAP forum, many of the companies involved decided to create a markup language (WML for wireless mark up language) that was closer to XML than to HTML.

In addition to micro browsers in mobile, you also have what is known as short message service technology. It bears similarities to two-way paging. People are continually connecting it to the Internet and to email. An email from the Internet may be sent to a mobile phone number with the address format of xxx xxx xxxx@phoneco.net where xxx xxx xxxx is the individual's mobile number and phoneco.net is the Internet address of the phone company. The first 250 or so characters of the message will be delivered to the recipient's mobile phone.

Today these capabilities are being used most widely in Europe on GSM systems which had been developed earlier than equivalent technology in the US. Many systems that allow messages to be sent to phones are now in the process of being up graded to allow two way messaging which is to say that phones will be able to send email back to Internet addresses. Because this technology is so simple, you cannot point to any manufacturers as leading the way. It is really built into all the standards.

**COOK Report:** As a result anyone with an Internet account can use his email client to

initiate a text-based paging call?

**Brodsky:** Yes. All sorts of interesting new applications are possible such as configuring the calendar on your desktop computer to send you a reminder message an hour before you are supposed to be able to begin an off site meeting? Or for example an enterprise can keep in touch with 500 sales reps all over the country who may all be traveling at various field locations at any given point in time.

Another application to which people are paying insufficient attention is that mobile technology is becoming robust enough so that in the next couple of years it will be possible to provide fairly high speed access to notebook computers and to PDAs. This will mean that a business traveler sitting at an airport gate waiting to board his flight will be able to up load and download email without plugging into an RJ 11 jack and using the PSTN. This kind of 64 kilobit per second mobile data service has already started in Korea in December 1999 and will be starting in Japan in January 2000.

These services are being provided by the CDMA industry which was started later than TDMA and GSM. CDMA now surpasses anything that TDMA can offer in 64 kilobit data capabilities. GSM is also planning to offer 64 kilobits. However, they are likely to have a difficult time making it available on a widespread basis because they don't have enough capacity.

**COOK Report:** And what companies are leading these efforts?

**Brodsky:** Qualcomm is driving CDMA technology forward but here are lots of other players. Motorola is the vendor that is actually launching the 64 kilobit service in Japan. Motorola had been late both in getting into digital and into CDMA and about a year ago finally got their act together on CDMA. They are now delivering CDMA phones and have become a pretty significant supplier of CDMA infrastructure. Motorola's financial performance in the cellular area improved dramatically in 1999.

LG, Hundai and Samsung are Korean companies making CDMA equipment. Japanese companies are also getting involved. Kyocera just bought Qualcomm's handset business. Fujitsu and NEC are also players.

**COOK Report:** But how do these networks transmit?

**Brodsky:** By means of CDMA at low power on regular cellular network channels. You can therefore send and receive data while you are moving. Note that CDMA is the only cellular technology that allows its users to talk to more than one base station at a

time. With GSM for example they have to shut down your call on cell site "A" and bring it up on cell site "B" at precisely the right moment so that your call is not dropped. With CDMA you could be talking to cell sites A, B, and C simultaneously.

## Metricom versus Sprint PCS

**COOK Report:** So where does Metricom fit in?

**Brodsky:** Metricom has a unique architecture that makes it a hybrid mixing features both of mobile and of the wireless LAN. Metricom uses frequency hopping spread spectrum.

**COOK Report:** Their architecture of transmitters hung from light poles spaced evenly across grids throughout an entire city has been around since 1993. The fact that they are getting 600 million dollars for a national build out indicates that someone thinks they don't have many liabilities.

**Brodsky:** They have done two significant things. First they brought in a new management team that is running with a tighter more realistic business model. They were functioning as a vertically integrated company where they were trying to be both the service provider and the ISP and the company providing the high speed modems.

**COOK Report:** In other words they were in the Internet business but saw themselves as a vertically integrated telco?

**Brodsky:** Yes. The new management team defined Metricom's core competency as providing high speed data access. They will work with many ISPs and will license their modem technology to multiple manufacturers.

**COOK Report:** So they will have a wireless infrastructure that anyone can plug into? They finally understood.

**Brodsky:** Correct. They did indeed. But in the longer term I do have a concern about their ability to compete in the new cellular/PCS world, because these people have a huge cellular PCS infrastructure to leverage in providing these services. When Sprint rolled out its wireless web service on the day of its launch, it was available from more than 11,000 base stations nationwide.

**COOK Report:** What can you do on Sprint PCS that you cannot do on Metricom or vice versa?

**Brodsky:** With Sprint it is coverage. It is available everywhere and the platform can be a phone that can accommodate both voice

and data. The key benefits of Metricom is that today it has a speed advantage because Sprint PCS runs at 14.4 kilobits per second where Metricom is getting its users absolute throughputs of 28.8. Metricom is rolling out technology now that offers 128k. However, the cellular/PCS guys will be rolling out 144 kilobits in 2001 and probably within the year after that will be getting up to 384 kilobit per second speeds. So the question for Metricom is what can they do to keep up with this? I am not convinced that they can. But they say that they have a plan.

**COOK Report:** In mobile data then is PCS in a position to become the dominant technology?

**Brodsky:** It is in a position to dominate. PCS is not so much a technology as it is a new category of players who are using the 1.9 gigahertz spectrum as their territory. All the PCS players started out from day one using digital and as such they have an advantage over the cellular carriers. The situation here is reminiscent of the situation faced by the incumbent telcos when they face competition from the new greenfield players like Qwest, Level3, Global Crossing and Williams. The incumbents will claim that they will have and offer all the latest technology. The problem for them is when they go to carry out their pledge to implement it, they are confronted with a huge and archaic installed base from which they have to migrate in a very measured fashion. This problem includes the very expensive necessity of figuring out how to get their analogue users to convert to digital.

**COOK Report:** Besides Sprint who are the PCS players?

**Brodsky:** Sprint is actually the largest PCS operator. The December 27<sup>th</sup> issue of Radio Communications Report (Vol. 18, No. 52) has a listing of the largest PCS players. Ranked by number of subscribers, they list Nextel the second largest after Sprint. Nextel, in my opinion is something of a special case because it operates in what is called the specialized mobile radio band which was created for business dispatch purposes. They got the FCC to relax the regulations that kept them boxed in to being purely a dispatch company. As a result they have begun to move out into a broader market where in my opinion they are not really competing with Sprint PCS for the general consumer but are competing with Sprint for the business user. What Nextel is doing is providing what they call workgroup communications. They are also pushing very heavily into Internet access. Primeco Personal Communications is third, PacTel Wireless is fourth, ATT Wireless is fifth.

**COOK Report:** Who else besides Sprint is

doing things with Web phones?

**Brodsky:** Nextel and then there are two European carriers. One in Finland the other in France. Air Touch and Bell Atlantic have also introduced them.

Air Touch, which is being acquired by Vodafone, is an interesting player. Air Touch was originally the Pactel wireless group. About three years ago they were spun off from Pactel saying that having to be part of an RBOC culture would not allow them to be sufficiently innovative. After they became independent, PacTel created a new wireless company which is a purely PCS operation. So Air Touch is actually PacTel's old cellular group. They were among the very first carriers to roll out CDMA. Slightly more than 60% of the PCS carriers use CDMA. There rest are split between TDMA and GSM.

**COOK Report:** So where does GSM fit in?

**Brodsky:** The Europeans were the first to make a wholesale transition to digital and in doing so they became some of the first to exploit the benefits of digital. GSM has grown to be huge. It is mainly used in Europe, but it is also used many other countries - many other countries have standardized on it. GSM is just a basic TDMA system on which many people have been working for years. It has very good international roaming which is very important to Europeans. Prior to GSM they had only a hodgepodge of analogue systems and couldn't use the same phone across borders, because the systems in neighboring countries were just a little different or operated on different frequencies. With GSM they created an international system with very good automatic routing features. Wherever you are when someone calls your GSM number, the network will find you.

**COOK Report:** With mobile TDMA as a part of GSM how does GSM handle the bandwidth issues?

**Brodsky:** Everything is relative. TDMA provided a 250% increase in bandwidth and speed over analogue. Therefore it instantly provided more capacity. The weakness of GSM is that compared to CDMA its capacity is limited. Now that digital wireless in Europe is really taking off, GSM networks in Europe are running close to capacity which is to say that the possible usage in a given area is just about filled. Finally, the audio quality that you get with GSM is not as good as with CDMA.

**COOK Report:** What can they do to increase capacity?

**Brodsky;** They can go to higher transmission rates but are somewhat limited there by

the risk of increased multi path interference. But they can also reuse frequencies more aggressively and apply new smart antenna technologies. They can also install more base stations and make the cells smaller. But the drawback here is that they have to do more cell handoffs - something that with increased frequency is difficult to manage well.

**COOK Report:** So both CDMA and GSM are expandable but presumably the cost per user of expanding GSM would be significantly greater?

**Brodsky:** Yes. Not only the cost but also the complexity.

## Wireless LANs

Now in the wireless LAN area the basic key is that companies have continued to achieve size, power consumption, and cost reductions. Providing momentum to these efforts is the realization that homes and small offices are going to need higher speeds. For several years wireless LANs were focusing on medium speeds as a way of getting lower power consumption and cost reduction. But now with further progress in making the digital technology, it is possible to go to higher speeds and to lower the cost. A few years ago a wireless LAN meant a five hundred dollar radio on a desk top and actually having to run a bit slower than a wired LAN. Looking out a year from now we are talking about a cost of less than \$100 per node and running at Ethernet speeds. Also these nodes can be portable within a building or even on a campus as well as fixed.

We are seeing a lot of activity here in terms of vertical applications. For example wireless bar code scanners have become very common in major department stores and in merchants like K-mart and Wall Mart. Hospitals are using them. In a few places on line shopping is even putting them to use via sales people equipped with notebook computers and web cams who can find merchandise and show it remotely to the customer prior to the purchase.

Another thing that becomes possible with a wireless LAN is linking home appliances together. Companies like Sharewave are working on embedding wireless LAN technology in a wide range of products. Now it is really the Internet that provides the glue to make sense out of all this. In the past to have your computer talk to your hi-fi and your refrigerator was pointless. Now with the Internet you may want to download music to your computer or to your hand held MPEG player.

**COOK Report:** You could have a web based interface that would allow you to remotely turn things on or off inside your house.

**Brodsky:** If you are having a problem with you refrigerator and your refrigerator can communicate wirelessly with your Internet connected home computer, then there's no reason why the manufacturer could not use that link to run diagnostics. I think that we are heading toward the cost of \$20 for putting a transceiver inside an appliance and using the device to connect the appliance to a wireless LAN. All sorts of interesting things flow from this. For example it means that big consumer electronics companies can create network out of all their existing products. You will soon be running your stereo, appliances, thermostat with the remote that was formally good only for running your TV.

**COOK Report:** Of course then to give IP addresses to everything becomes such a task that the use of IPv6 appears mandatory, But I am beginning to wonder if some of these technology developments may be running into some real problems in areas like routing due to the compromises in the deployment of routing architectures undertaken to scale the network in the mid 1990s.

For as network topology and its changes grow more and more pronounced, a network routing guru like Sean Doran points out that scaling on the edges of the network becomes increasingly difficult as broadband deployment moves into every home. While proxies involved in IPv4 NAT boxes can be traded off for IPv6 addresses aggregating huge amounts of bandwidth at the edges into smooth backbone flows could become difficult. This is a related area of complexity of which application advocates seem to be unaware.

**Brodsky:** you raise interesting points. Meanwhile I'd like to point out two other areas of application development. One is the push towards the development of ten meg or faster wireless LANs that will run throughout a home or small office. You have a battle going on here between the proponents of frequency hopping and direct sequence. The frequency hopping camp is lead by Proxim and the direct sequence camp is lead by Harris' Intersil and Lucent Technologies' Wavelan product.

**COOK Report:** What about Blue Tooth?

**Brodsky:** Blue Tooth is a low power system that is meant more for device to device communications rather than device to LAN communications. Blue tooth would allow you to have your cell phone talk to your computer while you moved about the room. In many respects Blue Tooth has the capabilities of a full blown wireless LAN with out the power amplifier need to extend the ranges involved to a couple hundred feet.

So with a 100 miliwatt power amplifier it becomes a technology capable of compet-

ing with a full fledged wireless LAN or a technology that serves as a bridge — allowing both to be done. It could make it possible to download a catalogue of material on demand from a PDA to a notebook or vice versa. You walk over to the customer and beam the information over to him. You could do this with Infra red. But here you have to aim and shoot. With Blue Tooth there is no need to aim because it can go around corners and even through walls for a short distance.

**COOK Report:** How would you characterize all this? In the wired world you have huge developments going on — this includes vast amounts of new applications bandwidth and involves lots of new players. You have lots of new technology getting cheaper and cheaper.

But your point is that one is missing the totality of the growth picture and the totality of the complexity of the architectural picture of global internet communications unless you also keep an eye on these wireless developments and have some level of understanding of how their existence is going to affect the net's wired players.

**Brodsky:** Yes and there are definitely two levels on which this is happening. One is that wireless is the key to really fulfilling the vision of cyberspace and that the idea of Cyberspace not that it would only be available in a desk top computer in a specific location where there happened to be some sort of wired connection. The idea of cyberspace is that ultimately the network would be everywhere and wireless is the key to making that happen.

Then the other aspect is that wired solutions have limitations of their own. Wired requires human labor to install and maintain. It also has an element of delay in time being necessary to provision service. You build the network first. Then you test it and only then can you allow service.

Wireless allows you to provide service though out an entire coverage radius simultaneously at the flick of a button just by putting a base station up. If the coverage radius is a mile when you have the base station up, everyone within a mile of that base station is essentially on the network.

**COOK Report:** In other words your infrastructure implementation simply turns on with the flick of a switch.

**Brodsky:** When you look at the really big picture comparing wireless and wired long term, the big advantage is that wireless can actually get us where we are trying to go which is bandwidth on demand. Get whatever bandwidth is needed to whomever needs it and do so when they need it. People

are always going to want more bandwidth. You will never be able to provide everybody with infinite bandwidth but wireless is the one to provide bandwidth on demand. If I need 100 megabits per second, to satisfy my need by wire will mean some constraints. COAX or fiber and probably just fiber is your point of delivery. But if I am wireless and only need 100 megabits per second one hour a week, it becomes much easier to provide the 100 megabits by wireless for that one hour.

**COOK Report:** But no one is looking for wireless to replace OC192 backbone links are they? These will remain all fiber will they not?

**Brodsky:** I agree. But, except for the very biggest links, wireless is becoming practical for more and more of the network. I expect that a year from now we will begin to see wireless replace dial up modems because wireless will actually be faster.

## Interoperability

**COOK Report:** What interoperability issues exist?

**Brodsky:** Analogue mobile phones had some issues but when you go digital, doing so meant that different approaches to operation were largely contained in software. Consequently it becomes relatively easy to build phones that can conform to multiple standards. Today you have phones that can do both analogue and CDMA or analogue and TDMA. Since analogue is nationwide, users are assured that their phone will work anywhere in the nation. In the future you will see phones that will do both CDMA and TDMA or CDMA and GSM. It is somewhat ironic that when the dual standard phone were first released some of them had problems with their analogue reception capabilities because digital circuitry was being used to emulate analogue behavior.

**COOK Report:** What has to happen for an analogue carrier to be able to support digital?

**Brodsky:** Such a carrier will have to overlay its entire network by installing a digital base station with every cell site. From a purely defensive point of view, most of the cellular carriers have had to do this because the use of PCS has grown so rapidly that it has forced them to speed conversions that were taking place very slowly.

An analogue phone cannot of course pick up a digital transmission but most cellular carriers, except in rural areas, are now operating dual mode networks. Once a digital call is picked up by a base station, it becomes just a bit stream that is injected into the PSTN and can go anywhere.

## A Wireless Equipment Market Rather than Service Market?

**COOK Report:** Sources are saying that Cisco will soon introduce a full line of wireless gear to assist ISPs in establishing wireless connections to backbone providers and to their customers. What is more Cisco claims that the equipment will be offered at very cheap prices that will be well within the reach of small ISPs. This sounds like a possible major change in Cisco strategy. For 1999 did not go well for Cisco in terms of its gaining a foothold in the multi terabit backbone router market. Juniper is out with a router that is said to outperform anything that Cisco can offer. Avici and several other companies are expected to debut products at least as good. One has to wonder if Cisco is seeking new hardware for sale not to the 200 largest backbones of the world but to the approximately 11,000 ISPs that connect to these providers. One wonders whether this is a believable scenario for Cisco and whether it might be prepared to spearhead a market designed to by pass ILEC infrastructure.

**Brodsky:** you are talking about the confluence of wireless infrastructure and IP. There is a real opportunity here. When you look at the landline world, you have the problem of such a huge installed base of analogue equipment that it take entirely new players to build new infrastructure to work around these giants. However in the wireless realm, everyone is moving so quickly to digital and things are already digital at the ends, it is now going to be a very interesting endeavor to begin to migrate to all IP wireless operations.

CDMA in particular is already a packet switched based protocol. This means that there is an opportunity to create for the first time real end-to-end IP based wireless networks. You have the next generation telcos like Qwest which are busy serving big business customers. The wireless carriers can begin to use end to end IP to serve everyone. Under these conditions you could support, it seems to me, a business model where a company like Cisco could begin to focus on selling very good quality equipment to enable users to link into the next generation telco network backbone.

But when thinking about this don't forget that the ISP is still going to need access to the radio spectrum. There are two options here. Licensed and unlicensed spectrum. If Cisco is going to go out and sell to the ISP, very likely the only viable option is the unlicensed spectrum. Otherwise the ISP will need an alliance with a big player that owns

# The Disruptive Internet: Triumph or Chaos?

## Accurate Assessment of State of Internet 2000 Depends on Mix of Technology, Governance Efforts, & Network Engineering Issues

At the beginning of the new millennium the Internet appears to be globally triumphant. Blindly embraced by politicians who do not understand it, it is touted as the engine that will power the economy of the new century. The lover's embrace by the rich and powerful has helped to make it a fad which is increasingly in control of the stock market. Internet companies are valued not out of sound understanding of their business models and cash flow dynamics but because they are internet companies. The Internet has emerged in the 1990s as perhaps the supreme disruptive technology in the sense popularized by Clayton Christensen in his 1997 book *The Innovator's Dilemma*.

Christensen found that disruptive technologies generally "under perform established products in mainstream markets," however they are generally cheaper and simpler to use than the sustaining technologies that they ultimately dethrone. Examined from a slightly different perspective, according to *Telegeography 2000*, "disruptive technologies give customers more than they currently need or are willing to pay for. This is one reason why such technologies are usually commercialized first by customers serving niche markets such as data networks or CLECs."

Because such technologies bring to market a such a radically different value proposition than was here-to-fore available, those practitioners of the older sustaining technologies find themselves simply unable to cope. To name just two examples central to the Internet's triumph as a disruptive network: in the last three years: we have gone from the established wisdom that state-of-the art telecom networks must be intelligent, to the realization that they must instead be stupid. Also we have traveled from the certainty of knowledge that publishing is an activity whereby elite gate keepers package knowledge for the masses to a world where every computer owner can become a publisher. In this new disruptive, Internet-driven world the old realities of our once comfortable industrial world are turned inside out. The only way to cope is to hop on the bandwagon and hype processes you simply don't understand. Jeff Bezos, the CEO of Amazon.com, a company that has yet to report a profit, becomes *Time* magazine's "Man of the Year."

This report will give an in depth look at reasons for the success of the disruptive Internet. It also steps back far enough to grasp some critical, but as yet unresolved, contradictions

that will influence how an ongoing struggle for control of telecommunications, commerce and access to the world's knowledge will evolve. We have technology analysts, e-commerce analysts and legal experts bringing their respective expertise to bear as they tell their separate stories of who is winning and who is losing. What we lack is a look at the larger picture of how the interaction of these areas will determine success and failure over the longer term.

This report will also examine in depth the critical areas of technology change. As it does so, it will show how the efforts of pre-internet telecom entities to channel and shape both legal and political control of this technology no longer fit the conditions under which these entities must work. The processes developed to cope with pre-internet, circuit switched, telecommunications systems simply do not afford policy makers any way to deal rationally with Internet disruptions.

From the FCC to ICANN, policy makers found in 1999 that processes designed to afford rational decisions about the use of the technology can no longer keep pace with the speed of technology changes. Education about what is happening can help. But there are no longer appropriate pigeonholes and regulatory boxes that can sustain the old way of doing things. Democracy and technology no longer mix. If technology is to be consistent with public values, the public must understand the technology and be able to move from that understanding to a grasp of the stakes.

But understanding the stakes is difficult. Technology itself is increasingly complex. Point A does not always lead to Point B, even if Event C occurs. The impact of a new technology on its associated environment is increasingly difficult to judge. Consequently we are likely to find that many of the assumptions on the basis of which policy has been developed over decades are no longer valid. Ultimately, as the system gets sufficiently complex with fewer and fewer people able to direct its course, some may ask when does chaos set in and take control? Technology that our political leaders do not understand is now driving the formation of public policy. As such values independent of the technology can no longer influence public policy. These are additional disruptive effects of the triumphant Internet.

Those who dare to assume the role of strat-

egist in the midst of these changes must understand that their successful stewardship will depend on their grasp not only of what Gilder calls the ascendant technologies. They must develop at least two other skill sets. One is an understanding that, as Lessig has shown so well, the success and visibility of the Internet has brought to it a level of attention where the legislatures and the regulators will impact it whether the rest of us like it or not. [For Lessig's Code: and Other Laws of Cyberspace see <http://cookreport.com/lessigbook.shtml>]

The other is a skill set that will help strategists to evaluate the stakes behind the questions of what will become the engineering agenda for the Internet. The approaches taken now to the implementation of IPv6 will impact the diversity of the infrastructure of the future Internet. Arcane but very important debates like these will determine whether issues of control over end-to-end uniformity and network transparency take priority over everything else. [Feb. 2000 *COOK Report*, pp. 11-24.] Will concerns arise that time spent toward these ends may mean that other more critical engineering concerns are ignored? A feared consequence is that the Internet's ability to continue to scale may be endangered.

## Technology Changes

The year just ended saw a continuation of the bandwidth revolution begun by the real take off of WDM technologies in 1997. New developments in DWDM (Dense Wave Division Multiplexing) and optical switching have further multiplied the bandwidth available from a global binge of fiber deployment. A revolutionary drop in the cost of data storage combined with increased network speeds is making it possible to deliver data across a wide area network more quickly and more cost effectively than across the bus on the motherboard of a single computer. These changes made the dream of the network becoming the computer - a dream that was first articulated earlier in the decade - seem likely to become true. They also made possible the rise of a new application service provider industry.

The year also saw the full flowering of new global telecom providers built on the inexpensive infrastructure made possible by the new technology. Qwest, Level3, Metromedia, Williams, Enron, Global Telesystems, Global Crossing, Next Link, Teleglobe, and Above Net to varying degrees

are all examples of these new telecom giants. While Teleglobe has been around for a number of years in a rather different form, and Williams is a reincarnation of an earlier venture, all of these new players with the exception of Metromedia have been profiled within the pages of the *COOK Report*. Questions of interconnection and peering are still critical. Equinix's neutral Internet Business Exchanges to be built globally under contract with Bechtel will be the most high profile model for fitting the new backbone players together and enabling cost effective interconnection.

An important new business involving the sale of surplus bandwidth at exchanges is beginning to emerge. We learned about RateXchange in a phone conversation with Ross Mayfield its CEO on December 30. In New York and Los Angeles Mayfield is running what he calls The Real-Time Bandwidth eXchange. He says that RTBX is a switch-based exchange that facilitates the entire transaction of bandwidth between interconnected carriers. Benefits include significantly lower transaction costs, anonymous trade, immediate delivery, and guarantee of payment. He seeks to create a win-win situation where an ISP with a temporary bandwidth surplus can sell bandwidth to an ISP with a shortage. His ultimate goal is to commoditize trading in bandwidth by means of the sale of contracts for future delivery.

When Web based interfaces for creating in almost-real-time single wavelength circuits between two points become common Mayfield and by then many other players intend to have infrastructure to make this possible. Given that the Internet for the past few years has been in a blind race to make enough bandwidth available to meet huge demands, this technology will be likely to give a welcome rationality to this marketplace. For the first time it should offer a means of efficiently and quickly bringing available bandwidth into balance with demand, and thus allow venture entrepreneurs to build new infrastructure where it will be used profitably

As this issue's lead article on the state of wireless shows, we have many new developments in wireless reaching their maturity and pulling the whole structure towards more ubiquitous always-on connectivity and applications. Wireless, becoming digital in 1999, became also broad band. Almost all the kinds, number and varieties of activities conducted via the wired Internet are now beginning to be carried out on the Internet in wireless form. This vast growth of wireless infrastructure has fed on itself to lead to a continued explosion of net use and net traffic. Large growth in cable and DSL connectivity in 1999 also contributed to the increase in traffic and in the number of always-on connections.

In 1999 the Internet reached a critical mass where every player in commerce and business had to have a presence. It became unthinkable not to have an Internet strategy. The growth in internet participation fed on itself and made possible business models which came to subsume every part of human economic activity. These changes mean that we are looking at a telecom world in which the old standards of evaluation fail us. The revolution is calling forth new standards and new criteria for understanding the forces propelling change.

## Technology Changes Divide Management into Irreconcilable Camps

The revolution has given us two opposed ways of thinking about and acting to achieve network organization. One may be described as the Bellhead intelligent network and the other as the Nethead stupid net work. (Some suggest that "Bellhead" should be referred to as "Telephone Head" and "Nethead" should be called instead "Internet Head.") In the first intelligence resides in expensive complex switches in the center of the telephone network linking stupid edge device telephones. In the second fat bit pipes linked by idiot savant routers connect intelligent computers on end user desks. The intelligent "Bellhead" network is generally organized by an extremely rigid hierarchy from the top down.

The "Nethead" architecture is organized by the creation of structure by cooperating autonomous groups. In addition to being important mental constructs, both architectures are now grounded in billions of dollars of competing infrastructures. These competing infrastructures are the foundations on which telecommunications, electronic commerce and individual access to vast knowledge will be determined in the new century. The older infrastructure (the Bellhead) is inherently more subject to vertical integration and organization. In the Bellhead environment, vertical, centralized, and rigidly-controlled hierarchies are used to enable communication, while the newer Nethead companies tend to be organized more often horizontally as companies whose market strength is in the delivery of one or more layers of the protocol stack.

This report proposes the outlines of a new business model structure for identifying and understanding Internet players who will prosper and ride the disruptive wave successfully. Growth areas in Nethead oriented companies are for specialists in bandwidth, interconnection, content hosting, application outsourcing, email. A player here is either an infrastructure-providing specialist, or one who coordinates and manages technology specialties horizontally to provide services

to end users. At the transport, infrastructure-providing, level you have the continued growth of new telecom greenfield players. These companies are taking advantage of new technology and new market conditions to grow from nothing into billion dollar plus companies in only two or three years although the long term judgment of the market is not yet in as to their appropriate valuation.

Opposed to this new disruptive Internet business model is the old solid and stable vertically integrated telecom model followed in the United States by the ILECs. ATT, with its attempt to buy control of the cable industry to have a foundation from which to deliver all telecom services, fits into the same category. MCI WorldCom with its attempt to become one of the biggest world wide integrated phone companies by acquiring Sprint also exhibits the telephone headed business strategy of a gigantic vertical integrator of dissimilar technologies and services.

## The Two Approaches to Slug it Out in 2000

We find that one battle of the year 2000 is to be between the horizontal versus vertically integrated business models. The horizontal model that is TCP/IP-based over gigabit or 10 gig Ethernet over fiber - a stupid network with smart peripheries where new low cost services can be cheaply and quickly inter connected like lego blocks.

In the disruptive Internet model you build and market horizontal collections of services at given layers of the protocol stack. Internet players tend not to be vertically integrated. Instead they are bandwidth and transport players like Williams, Enron, Global Telesystems or end user aggregations of services like MindSpring, Earthlink, or AOL that run over someone else's infrastructure.

Disruptive internet companies are nimble and quick compared to the older players like WorldCom that are still driven by empire building in search of elusive economies of scale. In a vertically integrated company, you may think you can cut costs by eliminating duplicative services, but instead you have a management nightmare of complex systems where glitches can mushroom into multiple week outages like MCI's frame relay collapse during the summer of 99. Or ATT as it attempts to build what it sees as the "network of the future" out of CATV policy and infrastructure. These players tend to see the Internet in terms of what has come before - in other words in terms of what is already familiar to them. What has gone before is regional feudal monopolies, based on exclusive land-based rights and exclusive land-based responsibilities overseen by public

regulators enforcing a social benefit/social cost analysis. And it well may be that this model may be with us longer than we would like. We should not think it can be so easily replaced, or that such replacement would have only good consequences.

Can the old vertically integrated intelligent network players compete? Or should they have some kind of protected quid-pro-quo status as the baseline common carrier? It will be much more difficult when wireless means cannibalization of their local loop income rather than new income. Look for them to engage in regulatory plays like ICANN.

The horizontally organized players can see the disruptions coming much more clearly, push the development envelope better, and achieve economies of scale by specializing in transport services, or storage services or application outsourcing. The vertical integration of the telephone company operators was what made the Network Access Points (NAPs) fail with WorldCom in the 1996 - 1997 period. (Some maintain that with a carrier running the NAP, it was then put in the counter productive position of competing with its own customers.)

In contrast horizontal aggregation of services will enable Equinix's neutral exchange points. Vertical integration is needed to make all the pieces of the intelligent network operate and to pay for relatively few and very expensive centralized switches. To make vertical integration work, one needs to control as much of the telecom environment in which one operates as possible. Even so some astute observers point out that, as the Internet industry continues to grow, even the horizontally oriented Internet companies are likely to make vertical market plays if they are sufficiently cash wealthy and driven to expand market share. We must strive for educational processes that will make clear to regulators and the public alike what economic and policy consequences will flow from the two different approaches to the market.

## Ignored Perils of Scalability and Architecture

The plug and play nature of internet interoperability means that a complex and specialized hierarchical structure can be built - one that offers market opportunities for many different companies of many different sizes. Although the over all network can be quite hierarchical, because there is room for many companies of many different sizes, management of each company can be quite flat and therefore rapidly respond to changing conditions. Yet some compromises have been made in network management that can spell trouble further down the road. To scale

network growth and cope with available IPv4 numbers, design decisions compromising end-to-end connectivity were made in the early 1990s. As the long IETF discussion published in the February 2000 *COOK Report* shows, there is a fallout from these decisions. The fallout is an incipient battle over how some of the intelligent edge devices of the network will communicate through the center to their intelligent counterparts at the opposite edges.

Internet interoperability encourages specialization and rapid technology development because the specialized pieces will fit together and work together. However the Internet is now set to begin paying penalties for its success. The triumph of IP in global telecom has made legacy companies believe it is critical to their futures to dominate the implementation of IPv6 in order to preserve an illusory and unattainable capability for uniform control of network operations

Let us summarize the argument we have just made. The Internet is having a disruptive impact in three areas of global importance. First is that of global commerce as trade of goods and services increasingly move on line and economic power is derived from one's positioning within and understanding of the new infrastructure.

Second is the reconstruction of the infrastructure itself in the provisioning of telecommunication services globally. As we have pointed out, the Internet is enabling new infrastructure that can provide all manner of telecommunication services at a tiny fraction of prior costs.

Third is the revolution in storage technologies that is re-enforcing the dropping costs of bandwidth. As a result, in many cases, ones computer will expand to become the network itself because internet connected data storage options, in some cases, will be cheaper than local storage options. The trend here may become the internet as the global repository for information and knowledge with all the economic and political power implied by the conditions of access to such material. While the internet has been designed to be decentralized and not subject to single points of failure or control, its very success in commerce, in more efficient network technology and in cheaper data storage will target it in the coming months as something that more and more forces will be determined to contain, own and control.

## Control Points: Ephemeral or Real?

One view of how services in the new telecom world will be delivered is determined by an architecture of an idealized Internet where TCP/IP will always be able to be sent from end to end. Because of a series of compromises involving

private address numbers, corporations now sit behind firewalls that demand the imposition of many translation devices. The Bellheads see these devices as obstacles to the imposition of a centrally controlled or controllable communication path like IPv6. The Netheads don't see any of the Bellhead's reasons for concern. They are happy to use multiple protocols and bridges over otherwise non interoperable network links. The Nethead's outlook focuses on extensibility while the Bellheads, through trying to influence the roll out of IPv6, focus on control.

ICANN which has garnered most of its support from bellheaded legacy organizations, has waged a generally successful campaign to sell itself to the press. The major press, having nothing other than ICANN's sound bytes from which to claim a clue, generally has no idea that anything is amiss. ICANN meanwhile is determined to error on the side of control concluding that it has the ability to enforce a uniform roll out for IP v6.

The two sides, one riding a disruptive technology and the other trying to extend the lifespan of the old sustaining technologies, sit in a hostile face off as they role out competing business plans for the 'converged' future of voice and data. Moreover we argue that there is evidence that they are fighting about control issues when they should be focusing on scaling and inter operability issues. We are very sensitive to Sean Doran's concerns about scaling of network architecture and routing to cope with a soon to arrive onslaught of broadband generated by the net's intelligent devices in homes and small businesses at the edge of the network. This is very little talked about. Therefore one must weigh whether the Internet technical community in doing battle over the presumed uniformity necessary for IPv6 is actually depriving itself of the time and effort needed to face more intractable problems of scaling, including protocol diversity.

Up to this point the Internet has scaled by a series of kludges. These are kludges that leave no one happy. While the Internet is not likely to suddenly fall off a cliff, if these issues are not attended to, performance is likely to get uglier and less reliable.

Thus we predict that what is otherwise an arcane dispute among network engineers is likely to become a basis for an operational strategy attuned to the structure, economics, and philosophy of the vertically integrated telcos of the majority of the planet. These conflicting issues meet head on in the distribution of IPv6 and over ICANN which is in a position to try to coordinate the implementation in a effort provide the uniformity that the Bellheads want and believe they need in order to stabilize their world.

We find that a vertically integrated market strategy is usually telco operated and at odds with the horizontally oriented and more cooperative internet market strategy. However in the current chaotic period of expansion and emphasis on building market share, these categories do not always hold. For the changes that we have chronicled for the past eight years are focused on building market share as well as

on creating successful newly disruptive businesses. When this happens a business like Cisco may mature to a point where the MBAs looking at the balance sheets and stock prices begin to drive things more than the engineers.

As one astute observer confided to us: “vertical integration has become the way that everyone is trying to rule the Internet economy, and I have seen very specific presentations from Cisco and others saying explicitly that this is what they are hoping to do. It’s the way to maximize shareholder value, but at the expense of an Internet that works well and that is user-friendly (by which I mean users that can innovate, not just click mice). Similarly, I believe that we’re past the nethead vs bellhead debate into something far scarier, which is where netheads and bellheads both realize that they’re going to have to integrate their networks. Moreover they may be committing to doing so even when they know that they don’t have time to agree about big chunks of the end result. In other words, while the PSTN may be getting dumber, intelligence is being injected into the Internet in some interesting and unexpected ways.”

Under such conditions amazingly powerful technology will almost certainly run into snafus originating with kludges in protocol design and IP numbering made in the first years of the take off of the commercial internet. When this happens more than web site performance will suffer. Also the profits of those companies which have marched forward fueled by nothing but blind faith in the continuing scalability and robustness of the Internet will take severe hits.

## Alternatives?

So are there alternatives? Ed Gerck who has participated in some of these debates has some wise suggestions in a long IETF discussion reprinted in the February 2000 *COOK Report*. Gerck advocates that the battle should not be thought of as either the uniform rollout of IPv6 or the failure of IPv6 and thus being fated to live with a network with some hosts blocked by firewalls or NAT boxes. There is a third way marked by diverse protocols that can inter operate.

When we asked him for further comment he replied “I saw no counter-argument raised to it — rather, it fell like an eye-opener to the reality that Nature is based on diversity. Those that believe in an uniform Internet as the only way to achieve end-to-end security are actually still locked in the network paradigm of the 70’s where network administration dictated orders to the entire network, by design. So, they will most naturally fight a multiple-protocol Internet, because they cannot intellectually cope with it.”

“However, in the Internet paradigm of the 90’s, now truly as networks of private networks in a progression over 30 years as catalogued by Stef, [Einar Steffierud] we already have multiple protocols in coexistence at various levels — what matters, so we have learned, is not that there must be one unique protocol at every place and time, but that different protocols, at different places, at different times, and doing different things are indeed able to work together when the objectives

are the same. And this, quite revealingly for the success of the Internet IMO, is how we humans prefer to work together, how commerce thrives and how we can enrich each other’s living experience.”

“So, communication protocols such as IPv6 must not be based on an excluding model, where they must kill any other protocol, but on an including model where interfaces are provided for backward and forward compatibility. Otherwise, even if IPv6 would be adopted by force on 90% of the Internet — this would still leave out 10% (which is projected at 40 million end-users in 2000) and would immediately raise questions about IPv7, IPv8 and so on. The only way to address the backward- and forward-looking questions without requiring the whole world to change standards at whatever cost at whatever time, is to provide for interoperation.”

“Thus, that is why I wrote that tools for IPv4/IPv6 interoperation will be needed .... and valued, including NATs as a fundamental building block — even though they began as a stumbling block in another context,” Gerck concluded.

Consequently, one must weigh whether the Internet technical community in doing battle over the presumed uniformity necessary for IPv6 is actually depriving itself of the time and effort needed to face more intractable problems of scaling, including protocol diversity. Those who would pretend to do large scale strategic planning in the midst of the internet revolution had better assimilate the lessons of Larry Lessig’s Code and Other Laws of Cyberspace.

It is time to realize that the Internet’s triumph as a disruptive technology now places it squarely at the center of attention of the lawyers, the politicians and the investment bankers. As the Internet moves into the legal and political arena the question for analysts to ponder is not a naïve belief in the unregulatability of the Internet. They need a much clearer understanding that the operation and impact of the Internet can be determined by the legal system, by the way in which its dominant protocols are coded, by the architectural environment in which it operates and by the customs that determine what behavior is socially acceptable.

Gerck’s statement about the need for interoperability tools for IPv4 and IPv6 is a good summary of why Lessig’s seemingly arcane issue of network architecture and control is politically very important. For in Lessig’s language, it offers in yet another way an example of how the technical decisions that change the shape of Internet architecture will indeed render it more subject to legal and regulatory control.

The disruptive Internet is in the midst of what looks to be a triumph. However the conditions affecting the outcome are far more complex than generally acknowledged when one steps back to consider how the totality of the process that we have just described falls outside the purview of more specialized analysts. For the first time technology is likely to be as much affected by policy decisions of network design and global regulation as by the older unregulated frontier mental-

ity typified by the mad rush to add users, scale the network, and increase market share. Winning players must take the complexity not only of the onslaught of new technology into account, but also make the right judgements about the current efforts impacting governance and control of standards in the face of escalating market importance but rapid shift in market economics. The winners will be those who get all the answers to this very complex mix of questions right.

---

## Wireless - continued from page 6

a block of spectrum. Now in the unlicensed area, one emerging model could look like Metricom where some companies become unlicensed spectrum providers for ISPs.

**COOK Report:** Then how feasible do you think spread spectrum is as an access mechanism - especially in urban areas?

**Brodsky:** I think that instances where you would experience unacceptable levels of interference in the use of spread spectrum by ISPs to deliver internet access in urban areas would be very rare. If you did have them it would be very nasty for the ISP whose access to the spectrum will not be protected.

Still I think that a couple of things have to be kept in mind. First Metricom actually did a nationwide study where it went out and listened to these unlicensed bands because their whole business model is at risk if these bands don’t work. They found out that even though there are all these consumer gadgets in the unlicensed bands, almost all activities are in the unlicensed spectrum at the 900 megahertz range. Now, in 2400 megahertz spectrum, the only things operational are microwave ovens that are designed to leak very little. Furthermore in the 5.8 gigahertz spectrum there is almost no activity.

Note also that, if spectrum does begin to become crowded, you will see money spent to counteract interference. For example you may have radios designed to search large areas of spectrum, find on a moments notices an empty spot, be able to communicate from there and then move on to some other slack area of spectrum when that gets noisy.

Also you must remember that in general spread spectrum has built in immunity to interference. For example back in the early days at an airforce base, the airforce used some unlicensed spread spectrum radios in the same spectrum that ordinary radios were operating in. None of the radio operators were ever aware of the others’ presence. What limitations exist are primarily on the regulatory side of the equation where if you ever do run into interference there is no one with the power to intercede on your behalf. You have to remember that you are limited to one watt of transmission power and that even hams are ahead of you in priority and can transmit with unlimited power. All in all with these minor caveats, I am quite optimistic about the ability of ISPs to consider a wireless access model especially if they have a powerhouse like Cisco making equipment for them.

# IETF Debates IPv6 Implementation and End-to-End Architectural Transparency

## NAT Boxes and Firewalls Seen by Some as Kludges to Be Eliminated and by others as Symbols of Healthy Diversity

Editor's Introduction: During 2000 the engineering of the Internet could be on the verge of falling victim to its own success. The successful development of DWDM coupled with a massive global build out of fiber has unleashed a flood of bandwidth along backbones. With DSL and cable modem deployment undergoing very rapid growth and with programs underway in Canada and Sweden designed to move fiber into residential areas, you have other pieces of Internet technology designed for earlier days of hierarchy faced with problems of scaling.

The explosion of Internet use has led to an increased demand for IP Numbers. This demand for numbers in turn lead to changes in IP block allocation. Classless Inter Domain Routing (CDIR) and Network Address Translation (NAT) boxes, both of which were designed to help ration the assignment of IPv4 numbers, were two results of these changes. CIDR made it possible to restrict the numbers of routes advertised across global backbones to the extent where the total has only doubled over the past five years.

In March of 1994 RFC 1597 Address Allocation for Private Internets made it possible to hide almost limitless numbers of IPv4 addresses behind Network Address Translation (NAT) devices. Because the IP numbers of these addresses never had to be advertised outside of the NAT device, they could be used again and again inside each different corporation. This however created major problems of new types. End to end transparency across the internet was lost. Protocols which depended on reaching a real human user at a unique IP number behind the NAT box would break. Because no one knew with certainty the number of IP addresses behind NAT Boxes, no one could say how severe a presumed worldwide shortage existed. Nat boxes, while a huge boon to many, broke end to end connectivity for many others and have made it difficult to estimate how many IP addresses are in use.

As became evident in a very rich discussion on the IETF mail list between November 26 and December 17, there are some very good engineers with very strong worries about the feasibility of an IPv6 rollout and network architectures and protocols that don't founder on the problems created by NAT boxes.

The discussion mentions version 4 of Brian Carpenter's "Internet Transparency" Dec 1999 draft. We quote from version 5. <http://www.ietf.org/internet-drafts/draft-carpen-ter-transparency-05.txt>. In his Introduction Carpenter states:

"The Internet is experiencing growing pains which are often referred to as "the end-to-end problem". This document attempts to analyze those growing pains by reviewing the current state of the network layer, especially its progressive loss of transparency. For the purposes of this document, "transparency" refers to the original Internet concept of a single universal logical addressing scheme, and the mechanisms by which packets may flow from source to destination essentially unaltered."

"The causes of this loss of transparency are partly artifacts of parsimonious allocation of the limited address space available to IPv4, and partly the result of broader issues resulting from the widespread use of TCP/IP technology by businesses and consumers. For example, network address translation is an artifact, but Intranets are not."

"Thus the way forward must recognize the fundamental changes in the usage of TCP/IP that are driving current Internet growth. In one scenario, a complete migration to IPv6 potentially allows the restoration of global address transparency, but without removing firewalls and proxies from the picture. At the other extreme, a total failure of IPv6 leads to complete fragmentation of the network layer, with global connectivity depending on endless patchwork."

**Editor:** Carpenter describes the importance of the development of corporate intranets. An intranet is "loosely defined as a private corporate network using TCP/IP technology, and connected to the Internet at large in a carefully controlled manner." He adds: "Network address translators (NATs) are an almost inevitable consequence of the existence of Intranets using private addresses yet needing to communicate with the Internet at large."

**Carpenter** continues: The notion that servers can use precious globally unique addresses from a small pool, because there will always be fewer servers than clients, may become anachronistic when most electrical

devices become network-manageable and thus become servers for a management protocol. Similarly, if every PC becomes a telephone (or the converse), capable of receiving unsolicited incoming calls, the lack of stable IP addresses for PCs will be an issue. Another impending paradigm shift is when domestic and small-office subscribers move from dial-up to always-on Internet connectivity, at which point transient address assignment from a pool becomes much less appropriate.

Many of the inventions described in the previous section lead to the datagram traffic between two hosts being directly or indirectly mediated by at least one other host. For example a client may depend on a DHCP server, a server may depend on a NAT, and any host may depend on a firewall. This violates the fate-sharing principle of [Saltzer] and introduces single points of failure. Worse, most of these points of failure require configuration data, yet another source of operational risk. The original notion that datagrams would find their way around failures, especially around failed routers, has been lost; indeed the overloading of border routers with additional functions has turned them into critical rather than redundant components, even for multi-homed sites.

The loss of address transparency has other negative effects. For example, large scale servers may use heuristics or even formal policies that assign different priorities to service for different clients, based on their addresses. As addresses lose their global meaning, this mechanism will fail. Similarly, any anti-spam or anti-spoofing techniques that rely on reverse DNS lookup of address values can be confused by translated addresses. (Uncoordinated renumbering can have similar disadvantages.)

The above issues are not academic. They add up to complexity in applications design, complexity in network configuration, complexity in security mechanisms, and complexity in network management. Specifically, they make fault diagnosis much harder, and by introducing more single points of failure, they make faults more likely to occur.

### Scenarios: successful migration to IPv6

In this scenario, IPv6 becomes fully imple-

mented on all hosts and routers, including the adaptation of middle ware, applications, and management systems. Since the address space then becomes big enough for all conceivable needs, address transparency can be restored. Transport-mode IPSEC can in principle deploy, given adequate security policy tools and a key infrastructure. However, it is widely believed that the Intranet/firewall model will certainly persist.

Note that it is a basic assumption of IPv6 that no artificial constraints will be placed on the supply of addresses, given that there are so many of them. Current practices by which some ISPs strongly limit the number of IPv4 addresses per client will have no reason to exist for IPv6. (However, addresses will still be assigned prudently, according to guidelines designed to favor hierarchical routing.)

Clearly this is in any case a very long term scenario, since it assumes that IPv4 has declined to the point where IPv6 is required for universal connectivity. Thus, a viable version of Scenario 5.3 is a prerequisite for Scenario 5.1.

### Complete failure of IPv6

In this scenario, IPv6 fails to reach any significant level of operational deployment, IPv4 addressing is the only available mechanism, and address transparency cannot be restored. IPSEC cannot be deployed globally in its current form. In the very long term, the pool of globally unique IPv4 addresses will be nearly totally allocated, and new addresses will generally not be available for any purpose.

It is unclear exactly what is likely to happen if the Internet continues to rely exclusively on IPv4, because in that eventuality a variety of schemes are likely to promulgated, with a view toward providing an acceptable evolutionary path for the network." [End of Carpenter transparency draft excerpts.]

**Editor:** for Carpenter with Intranets, firewalls and NAT boxes as the current reality, the only escape from lack of end-to-end security and the creation of a more easily manageable network that such end-to-end visions require is the implementation of IPv6. However, as we shall see from the discussion to come, the very existence of the many devices interfering with transparency will make IPv6 difficult to implement. One overarching question was posed by Ed Gerck toward the end of the debate when he said: "if NATs are to be blamed for the demise of IPv6, or its ad eternum delay, then maybe this is what the market wants — a multiple-protocol Internet, where tools for IPv4/IPv6 interoperation will be needed ... and valued." This seems to be a reasonable view of the world yet, presumably because it makes end

to end data security more difficult to attain, a good many engineers are violently opposed to it.

## The IETF Discussion

In the above general context, Eric **Flieschman**, [Boeing Corp.] on November 29, 1999 asked on the IETF mail list:

1) If we effectively ran out of addresses when RFC 1597 was published, has running out of addresses hurt us in any way? (2) Originally we had anticipated using IPv6 to save us from IPv4 address depletion. What's the status of IPv6? How does IPv6 traffic compare in volume with IPv4 traffic? Do non-IPv6-supporting vendors (e.g., Microsoft) have plans to eventually support IPv6? (3) Given the current situation of corporations using private addresses internally and a smaller set of global IPv4 addresses on their periphery, and a global IPv4 internet, one should theoretically be able to say how many public IPv4 addresses have been assigned and therefore how many remain unassigned and by so doing estimate how long until consumption. Why is this not possible in practice?

**Vladis Kletnikys:** Well, because it's a moving target.

If it takes 3 weeks to get the data together, and your growth rate is high enough that there's a measurable change over 3 weeks, you then have to apply a correction factor to your estimate. Given that you don't know the growth rate exactly, this gets more than a tad interesting.... Added in to that is the difficulty in computing how many IPv4 addresses are actually assigned - you can get a \*vast\* over-estimate by counting up the CIDR blocks visible to the core routers. However, in the case of Virginia Tech, you'll see the 128.173/16 and 198.82/16 net blocks advertised - but we're nowhere near 128K addresses actually allocated (I think the actual number is closer to 20K or so).

Most of the router gurus I've talked to don't think the actual number of addresses is an issue - the issue is allocating them such that a "small enough" number of CIDR aggregations covers them (as 30K poorly allocated IPv4 addresses and 300M well-allocated will both take 30K routing table entries.)

**Vernon Schryver** [in answer to Carpenter's IPv6 support question]: I used to work (and not as a manager, system administrator, salesperson, or other less technical job) for a UNIX vendor that had been working on IPv6 for a matter of years when I left. Shipping support for IPv6 in the same sense as shipping support for IPv4 is not as trivial as it sounds to managers, network administrators, salespeople, and standards committee

go-ers. It would be merely a major pain if all anyone cared about was sending IPv6 packets, but there are other things, such as naming, routing, and the reason for the exercise, applications. And no, while the changes to applications are generally conceptually trivial, in practice finding and making the required changes is not. If nothing else, the size of the bloated IPv6 address has painful, time consuming implications for applications, standards committee choruses not withstanding.

The number of public IPv4 addresses that have been assigned can be known; simply ask ARIN &co. However, I don't see how past assignments imply anything useful about the future. The vast majority of currently assigned IPv4 addresses could be re-assigned for varying costs. At one extreme, it would be almost free to reassign many of old class-C's that are not currently routed. In the middle, you could force corporations that own A's and B's to reassign to RFC 1597 numbers in their internal networks and yield their blocks. At the other extreme, you could force routed networks to be compressed.

In other words, the question is meaningless unless you also somehow specify how much pain you can or will impose or tolerate until the end and what you consider the end. Regardless of one's beliefs in the futures of IPv4 and IPv6 and whether you agree with Brian's apparent personal conclusion, draft-carpenter-transparency-04.txt lists many of the relevant issues.

**Keith Moore** [in answer to the question of whether running out of addresses has hurt]: Absolutely. Partially as a result of (perceived or actual) address scarcity, people have deployed NATs everywhere, and this has had an large adverse effect on the Internet's ability to support certain types of applications - particularly distributed applications (multi-party conferencing, games, simulations, distributed computations) and applications where the "user" end needs to be able to accept incoming traffic which is initiated from "outside". NATs also impair the reliability of the network because they can discard address mappings for connections that are still being used by applications.

**Tony Hain** [Microsoft] : (1) Yes [running out of addresses has hurt]... We have been forced into a world of NAT where simple applications such as Real Player won't work without real-time manual intervention at the NAT. (2) Yes Microsoft plans to support IPv6, and work has started. (3) We have moved from a world where Internal/External (and the associated management burden) was an option, to one where it is required. If corporations wanted to remove their firewalls (using IPSec instead) they couldn't.

**Steve Hultquist:** Many of the people who

have deployed NATs are responding directly to the address scarcity (and resultant cost). If you consider that many ISPs now have different pricing models for multiple IP addresses than they do for a single (regardless of bandwidth used), it isn't surprising.

**Moore:** No, not surprising (and in fact some people predicted it), but it is surely unfortunate.

**Hultquist:** I also think that it's interesting to consider that security concerns are the other primary reason for use of NAT.

**Moore:** This is indeed used as a selling point for NATs, but it's just snake oil. The security benefits of NAT are dubious at best.

**Richard Shockey:** You want to run a web server, DNS (from home?) Internet Print Protocol accessible printer, SIP phone .. pick your application. To the service providers the request for IP addresses is some sort of strange signal that you're running a eCommerce site from home or worse a Game/Porn site. Something that they believe they should charge more money for. The path of least resistance is just install a NAT. The market for NAT device/software products are being driven by in part the lack of IP4 numbers but certainly the lack of knowledge, marketing skills or just plain greed of the ISP community.

**Ian King:** Yes, my ISP is charging me for my DNS entry — a static entry made once in a zone file, but I'm paying monthly to have a domain name! Never mind that I don't use their mail server or Web page service or DNS server.... And yes, additional IP addresses were going to cost dramatically more. NAT was a simple case of economics... but on the other hand, I don't experience any "lack" because of it. I don't play UDP-based games or employ any of the other relatively new protocols that are so sensitive to end-to-end-ness (should they be? was that a valid assumption?), so a NAT is a great solution for me.

**Moore:** Understood. And you may never miss any of those distributed applications or applications that want your end to be the "server" for the very reason that NAT prevents them from having enough market share to be successful. That is to say: just because you don't know what you're missing doesn't mean that NAT hasn't done you harm.

**King:** NAT would be bad if an ISP were using it to artificially expand its address space; the use of NAT at the "small-time" end user's site seems quite practical and beneficial, especially in a world where ISPs are going to hold up non-naive users for ransom. **Moore:** If you think of NAT as a short-term strategy and are fully aware of its limitations, it probably won't cause you much

harm as an individual. Then again, there are dozens of products out there claiming to offer something like "internet connection sharing" without bothering to mention the limitations of that approach...which seems like misleading advertising at best. See <http://www.cs.utk.edu/~moore/what-nats-break.html> for my attempt to list the various ways that programs can lose in the presence of NATs.

On November 30 **Paul Ferguson** [Cisco] I sit behind a NAT/PAT device and Real Player works just fine for me. I've only found a couple of applications that will not work for me (e.g. ICQ, NTP, SNMP), but then again, I'm not a gamer so I can't speak to the broader range of applications that it \_does\_ break

**Senie:** Real added features to their protocol which permit it to work around most anything. These include using TCP instead of UDP if UDP doesn't work (probably how you're getting your streams, if you look at the statistics). I have found NTP (ok, SNTP) actually works fine through the NAPT implementation I use. A very large percentage of the protocols used by actual end users really do work, provided the servers are out on the public network.

**Ferguson:** I've always personally been of the opinion that if applications don't work in the face of NAT, then the applications themselves are functionally deficient and should be fixed.

**Senie:** Indeed. While some in the IETF have stomped their feet and railed about how bad NAT is architecturally (something I suspect most folks concede) they've somehow believed it would be possible to offer a better solution and get NAT eliminated before it's widely deployed. Reality: it's extremely widely deployed, and must be taken into consideration when designing protocols. Those, like Real, who find ways to make their protocols work with NAT clearly will do better than those who do not.

I've have thought I'd get a lot of feedback on the draft I've been working on in this area, [draft-ietf-nat-app-guide-02.txt](#), but that's not been the case. New protocols should, in my opinion, provide descriptions of how they work or don't work with NAT. If there is a reason why they aren't going to work (carriage of port or address information), a description of how to build an Application Layer Gateway (ALG) should be provided. We are at a crossroads where architectural purity has intersected operational reality.

## Network Address Translation

**Lloyd Wood:** There are a number of differ-

ent types of NAT - which the NAT working group has described. These break application assumptions to a greater or lesser extent. I don't believe applications are functionally deficient if they happen to require bidirectional NAT or twice NAT along with a DNS ALG to function properly; rather, it's the NAT implementation that's clearly deficient. <http://www.ietf.org/html.charters/nat-charter.html>

**Valdis Kletnieks:** Well.. Urm... TCP and UDP both assume that one endpoint is talking directly in real time to another endpoint. The NAT problems only start when the protocol carries IP address/port information (such as the FTP 'PORT' command), and the NAT isn't aware of that protocol's translation requirements (If you see \*this\* at offset 80 of \*that\* packet, smash it to read \*foobar\* instead).

**John Day** [BBN]: I would tend to agree. As I have said elsewhere, NATs in and of themselves do nothing wrong. They are doing things within the Internet/Network Layer that are perfectly legal. (In essence, they are treating the network address in much the same way that X.25, ATM, and MPLS treat their addresses.) The problem is that applications lacking an "application address space" are using the Network address space inappropriately. In a sense, we are making the same mistake the phone companies made when they kludged their route-dependent address space to be location-independent (first 800 numbers and then mobile). They have since fixed their problem.

**Kletnieks:** I'll grant FTP an exemption, it came well before NAT units became prevalent (Was there an FTP-over-NCP before The Great IP Deployment?). However, I do agree that anybody designing a protocol in the last 3-4 years \*should\* have designed it to be firewall and NAT friendly. (Yes, I know that can be difficult in practice. I guess that's today's "Welcome to Reality").

**Hain:** This is the kind of BS that keeps these debates running. NAT problems exist anytime a connection originates on the public side and there is not a preexisting clear mapping to the private side. I didn't pick on Real Player at random. My house is connected through a NATPT, and my wife couldn't get a video stream opened back to her machine. If I had pre-mapped those ports to her machine, then my son would not have been able to get them on his. If I pre-map them, all the bogus security arguments about NAT become invalid. Clearly NAT has allowed me to create an environment which is not sustainable, but at least I know enough to know what the problem is, the average guy on the street doesn't stand a chance.

Yes there are problems with protocols that carry addresses, but ignoring encrypted traf-

fic that really amounts to acquiring and synchronizing deployments of ALGs. In the early stages this doesn't sound hard, but will vendors be willing to add new ALGs to 3 year old NAT hardware? Will they create an update process that is easy enough for the average user? Will the average user be able to figure out which NAT needs updating, and what version it needs? Add the fact that people want to encrypt their traffic for privacy, and one wonders why so much effort is spent on this dead-end technology.

**Paul Ferguson** [Cisco]: I wouldn't be so quick to characterize NAT as a "dead-end" technology. Personally, I think NAT is just fine, but I'm a self-proclaimed cynic and also consider myself somewhat of a pragmatist. In any event, it works for me, but I could certainly be in the minority.

I think most of the hoopla surrounding NAT's revolve around engineering purism. And I agree that statements that assert that NAT's provide some sort of "security through obscurity" are complete red herrings. Having said that, I ask you: What do you foresee as a realistic IPv6 transition plan? Dual stacks? I don't see it happening, to tell you the truth. (Maybe this 6-in-4 stuff will actually help here.) The truth is that NAT's allow organizations to deploy machines in networks which otherwise would not have enough address space. To say that NAT's are unequivocally evil is unfair, methinks.

**Day:** Actually, there was and the PORT command existed as a kludge. The preferred approach for the data connection was intended to be a fixed offset from the telnet connection. The PORT command, then called SOCK, was inserted because BBN TIPs hardwired printers and such to certain sockets. Not exactly an example I would recommend following. The PORT command should have been retired decades ago.

Referring to Ferguson's [I've always personally been of the opinion that if applications don't work in the face of NAT, then the applications themselves are functionally deficient and should be fixed.], **Melinda Shore**, Nokia IP Telephony: I'm certainly not going to disagree with you about that, but H.323 does not work through NAT without extremely sophisticated stateful inspection/rewrite capabilities in the NAT, and it will not work, period, if the signaling streams are encrypted. For better or worse (and let's not get into that), there's a lot of H.323 out there and there's going to be much, much more over the coming years. RSIP isn't going to work cleanly with H.323, either (although there are some rather disgusting things you can do in the application about that).

We can "should" until the cows come home (Hey! They're home!), but there's a lot of

real software out there that is broken by network address translation.

**Henning Schulzrinne** [Columbia and Lucent]: Agreed. Any protocol that wants to have out-of-band control will have difficulties with existing NATs (without ALGs). Thus, by saying "let's design NAT-friendly protocols", we are effectively ruling out a whole class of designs and only allow protocols with outgoing TCP connections (and possibly UDP request-response protocols).

**Pyda Srisuresh:** I don't think, that is what is being said at all. If you could design applications that can work with NAT enroute, without needing an ALG; that would be great. But, if the applications do require an ALG enroute (as in the case of voice-over-IP which uses out-of-band call-control signalling), then the application designers should also consider what it takes to build an ALG enroute. This is an issue not just with NATs, but also with firewalls and perhaps security gateways.

**Henning Schulzrinne:** In the case of telephony, it would mean keeping a TCP connection open continuously to some outside server that would then use that TCP connection to send incoming calls behind the NAT.

**Pyda Srisuresh:** Are you referring to the TCP connection so you can retain the name-to-address mapping? I suspect, you need to do more than just that to permit RTP streams across NAT boundaries. In which case, you are better off monitoring the name-to-address mapping as part of the ALG, instead of relying on a TCP call-control session.

**Henning Schulzrinne:** Thus, this is not just a matter of existing software, but fundamentally restricting protocol design to a very narrow class of design choices, choices which are basically inappropriate for anything related to efficient multimedia carriage (real-time multimedia over TCP isn't exactly a great option).

**Pyda Srisuresh:** Please see my comments above. Protocol design will be deemed successful only as they are deployed. It does not harm to think of their impact on NATs, firewalls and security gateways - all of which are already deployed in a large number of locations.

**Hultquist:** While it is important to focus on building protocols that are as functional as possible in as many different environments as possible, I find the statement that protocols are "functionally deficient" that do not take NAT and firewalls into account to be misguided. The ultimate goal of a network, in my mind, is to create an invisible connection between process running in distributed systems regardless of their location or connectivity. While protocol development is an

appropriate place to address issues introduced by lower-level elements of the overall system, development of the lower levels should be focused on making development at higher levels as straight-forward as is practical.

As has been discussed more exhaustively and ably by others than I am able to, NAT breaks this model. By introducing a single point-of-failure into the overall system and by also introducing artificial limitations linked directly to the temporary scarcity of address space, it is an anomaly in the overall development of the network system.

The overall design philosophy for the network system, at least in my way of thinking, is one of inclusion and direct communication. We should endeavor to develop with that mindset.

**Tony Dal Santo:** I really think trying to make everything "NAT aware" or "NAT friendly" is spending effort heading in the wrong direction. While I am forced to concede the wide deployment of NAT technology, I believe the problems it solves (or at least claims to) are better solved by other means, or require different solutions.

**John Day:** The problem is not to make applications "NAT aware" or "NAT friendly". The problem is to make applications "IP address unaware". What is an application doing exchanging and using names for things 2 layers below it? Sounds like a design for trouble if I ever heard of one.

**Tony Dal Santo:** Protocols are difficult enough to get right. I'd rather see time spent developing good algorithms than NAT compatibility. What about future NAT "features"? Are protocol designers supposed to guess at what new convolutions will be dreamt up?

**John Day:** You are absolutely right. Time should be spent developing "good algorithms" which is common "good architecture". What NAT does is just another form of the same thing that X.25, ATM, and MPLS do with different identifiers. It is not bad algorithm there nor bad architecture. (I presume since there are so many ATM and MPLS supporters (X.25 seems to be gone for the most part, thank god.))

**Tom Narten** [IBM, Raleigh, NC]: I find this assertion to be amazing. It's perfectly legal for a device to modify any field in the IP header if it so desires? Do you also agree that it's legal to modify the IP ident field (and potentially break fragmentation?) What about the length field, or flags field, etc., etc.??

**John Day:** Cmon, surely you can come up with a better counterargument than that! ;-)

I certainly could. If it is architecturally acceptable for those protocols to rewrite the address field at every hop, why shouldn't it be for IP? How does it differ? Basically a NAT is doing what an ATM switch does. How does an MPLS tag differ?

**Tony Dal Santo:** Here is the difference between NAT and the other things you mention. The only changes to the IP header and encapsulated data should be the TTL and fragmentation information. Granted ATM chops the packet into small cells, but when its put back together the packet looks the same. This cannot be said of NAT.

**John Day :** The answer of course is that MPLS tags, etc. have a much smaller scope and rely on the layer above to provide an address of broader scope "to bind the tags together." The problem for NATs is that IP doesn't have a name space of broader scope "to bind the tags together."

**Tom Narten:** Also, NATs have to modify the TCP/UDP checksum, i.e., look inside the upper layer protocol it is carrying and modify the payload. This is also not "bad architecture"?

**John Day :** That is more questionable given the architecture that has been put forward. If IP and TCP are supposed to be in different layers then the pseudo-header is also bad architecture. If they aren't in different layers, then it isn't and neither is modifying the checksum. [...] However there are so many special cases now of people doing strange things with IP addresses that they shouldn't be doing that there may not be anyway out of the problem.

**Tony Dal Santo:** Ok, here is where the rubber meets the road. What is "the problem"? My guess is that you think the protocols/applications are broke, or architecturally unsound. I think the problem is that NAT became "necessary".

**Narten:** If you want to point fingers, TCP is also broken from the perspective of NAT, because transport layer addresses are formed from both IP addresses and port numbers. It's not just applications sending addresses in payloads that "are broken", it's a key aspect of the basic TCP (and UDP) design.

**John Day :** This can be dealt with in the NAT.

**Narten:** Sure, there are folks that say "fix TCP". But this is also naive. There is a good reason why TCP is designed this way. Having transport addresses be completely independent of the IP address in the IP header requires having some way of mapping from one name space (e.g., the TCP transport identifier of a service one wants to communicate with) to the IP address (so that the TCP

header can be put into an IP packet and sent to the actual destination).

**John Day :** Not really. I recently heard Dave Clark say (and in a rare case I actually agree with him based on my own investigations into protocol design principles) that it would probably have been better if socket ids were part of the IP address and there was no addressing in TCP at all.

**Narten:** Although folks have talked about doing something like this for years, it has not been done (i.e., where is a document showing how it can be done?) and folks have argued endlessly about whether building such a system is feasible or really solves the problem as opposed to just creating a new (hard) problem elsewhere that then needs solving.

**John Day** Yes, I am aware of this. In the very early days, it was assumed that it would have to be done. Then there was a long hiatus where the only applications were telnet and ftp and by the time it was really needed it would seem that reticence to change had set in and lots of people continued to do ad hoc things that made change difficult.

I am well aware of the issues, but I did want to push back on people who see no problem with applications exchanging IP-addresses. The longer we keep putting this off the messier the problem gets. Frankly, I have heard these arguments for years that to fix anything will only create problem elsewhere and so far I have seen no examples that were not either "bad architecture" or should be accomplished in a different way. Where the "desired" approach was an artifact of not having the right tools to solve the problem.

**Tony Dal Santo:** Is there a constraint/expectation that these applications/protocols run on something other than IP? Is it a requirement that when (if) the switch is made to IPv6, that all IPv4 applications will work over IPv6, and even better yet work across an IPv4/IPv6 "converter"?

While these would all be nice things to have, are they design requirements? Is the goal to have a stackable streams-like system where I can slide in/out or replace "modules" letting me (in theory) run any application using any transport over any physical medium?

It's a nice idea, but certainly brings with it lots of complexity. My personal opinion is that the IETF's goals essentially boil down to "IP on everything". Unfortunately, this means if IP changes, everything much change with it. Just as IP's predecessor is gone, so would vanish old versions of IP.

This of course has its own problems, but I think identifies (at least for some) the current mindset. If the IETF has to identify and

specify the entire behavior from user interface to bit order on the wire, their charter needs to be expanded.

[Different reply of **DelSanto**] - About the only serious issues NAT really addresses are the perceived IP address shortage, and being forced to renumber when changing ISPs. As for the shortage, this thread seems to indicate there isn't a real firm idea of how many addresses are actually being used. Have all the minimally used class A and B addresses been reclaimed? I suspect not, which tells me the immediate shortage can't be that great. And if switching to IPv6 in the long term has problems, effort should be spent solving those.

As for being forced to renumber (can't take your address with you), is this really a technical routing problem (tables too large), or more of an economic/political/turf issue? While renumbering can be painful, DNS should hide the change, and things like DHCP can make it less so. Besides it shouldn't be too difficult for small organizations, and large ones should be able to take their address as they move.

**John Day:** Now you are on to the beginning of a good solution.

**Tony Dal Santo:** While NAT is an adequate stopgap solution to IP address dilemmas, in my opinion, it shouldn't be the final solution.

## NATs and Firewalls Break Protocols

**Ian King:** I think the above statement provides important perspective. NAT is not the Antechrist, nor is it salvation. Much of the work on "improving" NAT seems much like "improving" the Band-Aid so it will last for a year, although no one wears one for more than a couple of days! When IPv6 is deployed and everyone's toaster can have its own IP address, I suspect that most folks will be perfectly happy to decommission their NAT boxes.

Firewalls are another and likely more significant issue. However, focusing on firewalls narrows the issue considerably; how many corporations are concerned whether their firewalls are Quake-friendly? For those protocols that are of interest to users of firewalls, the necessary work can be done to either build ALGs, figure out tunneling methods, or design firewall-friendly protocols; that work will be driven by a business need, rather than an academic discussion of what "should" work.

It's important to know which protocols are broken by NAT and firewalls — Keith Moore's work on that is very useful. But does

each instance of “breakage” represent something that needs to be “fixed”? Part of this problem (NAT) will almost certainly go away; the other part (firewalls) requires at most a subset solution.

**Matt Holdrege** [NAT WG co-chair]: Bottom line, NAT or changing destination addresses on the fly breaks the end-to-end nature of Internetworking. This is true not only for IP, for other protocols as well. Some may recall older DECNET IV networks that exceeded the maximum ~65K nodes. We had to come up with a scheme to make that work and it was ugly.

Even though we are not in the same situation as DECNET IV (we still have a lot of addresses available), we are (for other reasons already discussed) stuck with this NAT situation for the foreseeable future. It does absolutely no good to complain about the state of the world. It is what it is and all the complaining by the IETF won't change it.

It would do a lot of good to find a way to dig ourselves out of this situation. The Internet isn't going to stop growing and there will be future applications that require or desire a heck of a lot of public IP addresses. Other future applications may compromise some useful feature to be “NAT Friendly”. This is not a good thing.

So we need to find a way to bridge ourselves to a new addressing architecture. IPv6 is one such architecture, but getting there is very, very difficult. So we need some alternative solutions in the meantime. And we need to continue to reassess how we might transition to such an architecture. The Internet will not stop changing so we need to be fluid in our transition schemes.

The IETF, IAB and IRTF are struggling with this and there is no neat short term fix. The NAT WG and the NGTRANS WG are a couple of places (I'm sure there are others) where folks are trying to help things along and we would appreciate your support.

P.S. Let me take this public moment to ask (beg!) for review of draft-ietf-nat-protocol-complications-01.txt which I'm editing. If you know of any particular protocol that has difficulty with basic Network Address Translation, please send me the info in the format of the above draft.

On December 1 **Noel Chiappa**: Everyone, this conversation isn't really going to be very productive. The people who like A aren't about to start liking B, and vice versa. (And then there are the people who don't like either - but they aren't going to change their minds either! :-)) So discussion on this point is not going to be very productive. The various approaches to growing the Internet (IPv6, NAT's, etc) all have costs and ben-

efits - but each differing costs and benefits. In a system as large as the Internet, it is to be expected that some people will evaluate the costs and benefits, and decide that for their particular situation, option A is to be preferred. Others will prefer option B.

**John Crowcroft**: Yes, but providers don't actually ASK the users what the COST is of a NAT. The BT ADSL trial in London uses NATs and all the folks I know who are what BT might call “advanced” users (i.e. people who want to export files on web servers, use multicast apps etc), find it a REAL PAIN in the backside compared to their old school/university ‘always on’ access but at no point in any of the market research on the trial users did the provider bother asking about this, so they remain clueless....as do many NAT Proponents.

**Noel Chiappa**: In the end, if a single approach does become the predominant one, it will be the market that decides (as individuals look at the costs and benefits in their individual cases), not the IETF. The IETF can't do a lot to influence this outcome - we don't have a police force, nor an army, so we can't stop someone from using A, or force them to use B. The IETF has always done much better at *designing* protocols than it has done at *choosing* protocols. Let's stick to designing, and stop arguing about which one to pick, eh?

On December 2, **Perry Metzger** [referring to Brian Carpenter's statement]: “My personal opinion is that we ran out of addresses in practical terms around about when RFC 1597 was published,” said himself “I hate to start a flame war, but Brian is absolutely right. I have clients spending gargantuan amounts of money dealing with layer upon layer of NATs because of this. Some people believe that “we aren't running out of address space” but the fact is that we've already run out. It just isn't available for many essential purposes.”

**Metzger** then added: “A client of mine just spent millions of dollars because of our current broken pseudo-internet in which everyone has to use private addressing 90% of the time. None of this money would have been spent if end to end addressing had been possible. This expenditure doesn't count the constant maintenance nightmares that layers of NATs have caused this particular client, and the expenses coming from that.”

“Money is a nice, quantifiable kind of pain. From what I can tell, the pain levels are very high. When your NAT nesting turns your simple network restructuring into a multi-million dollar game of pick-up-sticks, you're past the point where its ignorable. Of course, as one provider told me at the DC IETF, in literally these words, ‘NAT works just fine.

Besides, customers are mostly there just to spend money anyway.’ “

“When you've been awakened in the middle of the night every night for a week, because the NAT rules to deal with the fact that you have several intercommunicating networks all of which think they're 10.0.0.0/8 have become so complex no human can really remember them all any more, you'd experience what many of the rest of us feel week after week. The lack in question is a lack of sleep.”

“Anyone out there who thinks NAT works well and is harmless is not familiar with how the brave new world of pseudo-internetworking works. ‘Oh, no. When we moved the mail hub for the client access networks in location A, it seems no one remembered to update the NAT rules for the systems in location C. That's why the help desk has been inundated for three days...’ “

**Ian King**: “So what you're saying is that NAT doesn't scale well, at least from a management perspective. I'm not surprised. Again, I would suggest that NAT is not a silver bullet, but rather a useful “hack” in some circumstances. The situation you describe appears to be demonstrative of its limitations. (By the Way, one time I forgot to change the NAT rules when I moved my sendmail server on my four-machine network, and.... :-))”

“But then again, I would expect that a large corporation would see the need to own a large address space, rather than attempting to “pseudo-expand” its address space through the use of NAT. (My company, with a fairly substantial intranet, uses proxying as its internal solution.) Maybe I don't understand the usage you're describing, but the point I keep trying to make is that NAT isn't evil; that doesn't mean it solves every problem, OR that it should be re-engineered so that it does.”

**Metzger**: “You are assuming they could get such a space. They can't. No one can if you aren't a provider. The registries will simply tell you to go off in the corner and use private address space instead. The situation right now is really horrible. People are pretending it isn't, but the pretense really has to end. The End to End model isn't just a “nice idea” — its actually needed for people to manage and maintain networks cost effectively. Its effectively dead a the moment, and frankly, I'd like the internet back again, instead of this kludge we're using instead these days.”

**Yakov Rekhter** [Cisco] to Metzger's earlier remark: “Consider an alternative where the client decides to use IPv6. Granted, the client could get enough IPv6 addresses for all purposes, regardless of whether these

purposes essential or not. But then in order for that client to communicate with the rest of the folks, the client would likely use NAT (as the rest of the folks would still use IPv4).

On December 3 **Jon Crowcroft** : “its economics - if one large client uses IPv6, then so will others eventually as its cheaper for all of them collectively than for them to go on using NATs.”

**Rekhter**: “So, the cost of using NAT wouldn’t go away. But in addition, this alternative would cause the client to swallow the cost of transition from IPv4 to IPv6 in its infrastructure.”

**Crowcroft**: “right - the problem is getting the FIRST person to go - clearly a PROVIDER could consider swallowing the cost (i.e. pay Cisco to implement, and debug a deployed IPv6 backbone and then chase customer problems) - why? because in the LONG run there’s more subscribers selling content, doing business in a fully IP (v6) connected net than on a NAT disconnect, and so there’s more money for the provider...’ By the way, I think the address space stuff for subscribers using NATs is often (not always) hokum - it’s mostly that they can’t be bothered to design a decent addressing architecture for their intranets.”

**Bill Manning**: Oh, I think that there are lots of good engineers out there who do a great job of designing decent addressing architectures for their networks. Now if the darn things would just stay within the design constraints over time, there would not be a problem. Trouble is, those pesky users keep thinking up new things to do with the network, violating design assumptions. Makes the addressing architects life a living hell... (insert a quick overview of the evolution of the IPv4 addressing architecture here...) Still, I do think a basic NAT design philosophy is one of “withdraw into my own little world, protecting me from dealing with the troubles outside”.

**Noel Chiappa** [to Metzger’s complaint: “Anyone out there who thinks NAT works well and is harmless is not familiar. . .”] Perry, I’m curious about the technical aspects of the problems you’re seeing, in particular: Are the problems you are seeing due to i) the need for NAT boxes to grope around in packets, ii) the fact that hosts don’t have permanent, globally visible internetwork-level ‘names’, or iii) something else (e.g. complex configuration management)?

**Metzger**: All of the above.

i) hits all the time in that one is forced to tell people that some new service or another isn’t available because the various NATs and proxies don’t support it. That’s actually not

such a big problem, though, in the sense that if you don’t care about being able to use the net in new ways it doesn’t hurt.

ii) and iii) (complex configuration management) are the real issues.

**Chiappa**: The reason I ask is twofold. First, there is an alternative technology being proposed for local addresses in IPv4, RSIP, which should avoid i), but still leaves us with ii) and iii). So, to the extent the problems you are seeing are ii) and iii) we’re still kind of stuck, even if RSIP happens. To the extent that the problem is really i), though, RSIP might alleviate the situation.

Second, when examining the transition technologies for deployment of the proposed new inter network layer, I’ve been pondering the problem of an non-upgraded host trying to talk to an upgraded host with an address which is only expressible in the new inter network layer. The proposed transition technologies I’ve examined (i.e. NAT-PT, AIIH, and SIIT) \*all\* seem to have (at least as far as the IPv4 world is concerned) characteristic ii) - in that as far as non-upgraded hosts are concerned, upgraded hosts using those schemes don’t have a permanent ‘name’ at the inter network layer.

So I’d really be curious to know a little more about the “forest-level” nature of the problems you’re seeing out there - I think it will be very insightful in considering a number of potential forward directions.

**Metzger**: Having looked at things, RSIP appears to be as complicated as deploying IPv6 with local NATs to connect to the v4 internet, since it requires kernel hacks and border translators. v6 would have the advantage that at least after some years of deploying the hellish transitional technology the v6 bubbles would start to coalesce and form a real internet again, i.e. the pain would eventually end. RSIP would institutionalize the horror forever. For my money, I’d prefer v6.

By the way, I fully agree with those who contend that v6 does not solve the route aggregation problems we have in v4. However, people don’t seem to get that the raw address space size problem v6 solves is in and of itself reason enough to move to v6 given the costs we’re having trying to keep v4 on life support.

**Chiappa**: Perry, one other thing I’m curious about (based on your messages in particular, from this thread), and perhaps you can enlighten me.

While I’m not saying that NAT’s are a Great Thing, I do wonder if people are experiencing a bit of “grass is greener on the other side” syndrome here. NAT’s are in wide-spread use, and lots of people (e.g. you :-)

are being forced to struggle with their manifold downsides on a daily basis. In dealing with this (and wishing they didn’t have to deal with these problems :-), I wonder if people are assuming (perhaps incorrectly) that the alternative is relatively pain-free?

**Metzger**: If you mean “globally addressable IP addresses”, I think that the specific pain in question would go away, yes.

**Chiappa**: (My examination of some of the transition strategies indicates to me, as my previous email alluded to, that perhaps they are going to suffer from many of the same root causes that makes NAT painful too.)

So my question is: do you know of any site, one which has a lot of interaction with the rest of the Internet, which has actually converted to the alternative, for use on an ordinary daily basis, for the bulk of their activity? (This is, after all, the stress level we’re subjecting NAT boxes to, these days.) If so, it would be interesting to hear of their experience, so we can find out to what degree the results of a conversion are any “greener”.

**Metzger**: What do you mean by “the alternative”?

If you mean “gotten large blocks of global v4 addresses”, the answer is no one \*can\* at the moment. If you mean v6, obviously not — v6 isn’t in real commercial deployment yet (although it is now at the stage where we could commercially deploy). If people were to start deploying v6, it would probably initially be no less painful than NAT — it would only have the promise of being less painful down the road. If you mean RSIP, RSIP is even further from deployment than v6. Indeed, I’d say that RSIP is a clever but utterly dead end idea.

**Daniel Senie**: I too would rather see effort put into IPv6... if it’s going to happen, let’s get going, though... it’s been in the oven too long.

There’s one scenario about IPv6 which worries me, and which may well make RSIP and NAT in a pure IPv6 world a reality. Today we have LOTS of folks using NAT(NAPT really) for connecting to cable modems and DSL lines. Some folks are doing this for dialups too. It’s the model for “home networking” today. Will ISPs be willing to assign a block of addresses in the future to home networks? What does that mean when the access is a dialup?

Sooner or later, we’ll have providers handing out a single IPv6 address to any home user customer, just as is done today with IPv4 addresses. It is for this user population that RSIP will likely be a real issue. One of the things to think about is that while there is use of various NAT flavors in corporate en-

vironments, it is or will be nearly ubiquitous in the home market? Linux and Windows both include credible, functional NAPT solutions today. Is it ugly? Sure. Is there any chance it'll stop being attractive to the home user? Unlikely.

**John Stracke:** Getting people who have old non-aggregatable addresses to transition to v6 will give them the chance to get aggregatable addresses, won't it?

**Daniel Senie:** Perhaps. It then trades something else off. There are legitimate reasons to NOT have aggregatable addresses. Some companies actually like multihoming, as a way to keep their services operational in the face of one network provider melting. Aggregation isn't the goal of customers, reliability is. Having multiple connections to one ISP is not considered sufficient, either. We have adopted a world view where only ISPs are worthy of being multi-homed, and everyone else must aggregate through a single ISP.

There's something Perry may be alluding to (or not). Many companies create PRIVATE internetworking links among themselves for handling sensitive data. These require unique addressing. IPv6 has, to date, not been an option for this use. Two reasons: 1. The registries were not, until recently, handing out IPv6 addresses, and 2. The registries are using the same kinds of rules as for IPv4 in that addresses are handed out to ISPs who then hand them to companies. The problem with this is that if you never intended to route the addresses on the public net anyway, you won't be seen as worthy by the ISPs, or the registries.

I have been thinking for a while that it'd be useful for IANA to set aside a class A block for a new registry. This one would provide very small blocks of addresses, for private inter-corporate use only, to anyone willing to pay. This would provide a guarantee to users that the addresses they get ARE unique, and will work. It's kind of like RFC 1918 addresses, but with a block set aside for private interconnects. Clearly it's too late to do this with the blocks presently in RFC 1918, though. Take this same idea, and replace that class A with a big prefix in IPv6 space and perhaps the folks doing private interconnects will help fund the development of IPv6 equipment and updated applications.

**Metzger:** NAT has actually created a simple transition plan for us. I'd say at this point that 95% of the corporate networks in the U.S. use private addressing and a NAT or proxy box at the border. Switching from this to using v6 internally with a v6 to v4 NAT/proxy at the border for communicating with v4 is trivial — since they don't have globally routable addresses now, they won't be hurt by the switch.

As more and more people switch to this configuration, they'll start finding themselves talking to more and more things over the net natively, and fewer and fewer through the translator. Suddenly, they'll discover they \*do\* have globally routable addresses again, just like we did in the old days before net 10 was turned into the universal addressing ghetto.

**Rekhter** (on December 4): This could go as you described \*until\* these folks would start to move from one provider to another, and therefore will be faced with the need to renumber. At that point NAT could become quite an attractive alternative from cost/benefit point of view again. And now these folks will be back to where they've been with IPv4, but with the extra cost of IPv4 to IPv6 transition. While one could argue that renumbering with IPv6 could be made simpler than with IPv4, what one really needs to compare is renumbering with IPv6 vs renumbering with NAT.

To Noel Chiappa's question about the problems that Perry Metzger sees **Steven M. Polinsky**, Vice President, Information Technology, Goldman, Sachs & Co. added on December 6: To me the biggest problem here, is the common situation such that companies have separate (and necessary) Internet and Remote Access firewalls. RA firewalls exist in multiple global locations within an enterprise.

Multiple instances of the same Private addresses would enter (or exit) the enterprise network via Private lines from different companies if not for careful configuration management across and negotiation between "NAT Administrators", within the enterprise, and between enterprises. The most difficult part is the negotiation with client/vendor site NAT Admins as to who should NAT which addresses into which addresses. We often need to negotiate between 3 RA connected companies. Not only is this painful, but one can never sleep comfortably, knowing that a NAT Admin at a 3rd company will not make a mistake and connect someone new at our NATed address. There are not enough Private Addresses to go around. I would propose that an additional block of addresses be added to those designated in RFC1918. That would be a big practical help.

**Kieth Moore:** the only way you can have private address space large enough to map all of the other private address spaces (that you wish to talk to) into your private address space, is to reserve a portion of your own private address space that is large enough for all of the other private address spaces (that you wish to talk to) combined. problem is, everyone else's private address space is potentially as large as your own. making the private address space larger won't solve that problem. it might give you

a few months' breathing room, but eventually you end up having to limit the amount of address space that any organization can use, so that there's room to map it into everyone else's address space. once you do that you might as well just assign a chunk of address space to each organization. and for that we need IPv6.

**Randy Bush** (on December 7): a LOT of folk have deployed NAT. Hundreds every day. It's easy. It solves the customer's perception of their problem. It's not expensive. [Yes it sucks architecturally. What's that got to do with customer/market reality? The office buildings we work in suck architecturally too.]

We'll have to be at least three times (cheaper\*better\*easier) to replace NAT, that's market reality. Currently we are not proposing anything that the customer perceives as cheaper, better, or easier. This does not bode well. Until we change this by making our technology, or at least the perception of it, cheaper, better, and easier, pontification, scare tactics, and nat-bashing make us look foolish and poison our credibility in the long term.

**Carpenter:** The idea is that IPv6 site renumbering will be so much easier than for IPv4 that renumbering will be \*less\* painful than NATing.

**Bush:** This needs to be reconciled with the \*much\* more conservative statements on v6 renumber-ability coming from respected v6 folk such as Deering et alia.

**Metzger** (to Bush's statement that NAT deployment is not expensive): It is \*astoundingly\* expensive. It only seems cheap until you have to maintain it. And yes, I'm going by Actual Live Customer Experience In Actual Live Large Companies.

I'll keep posting this so long as people keep on with the "NAT is cheap and works well" myth. NAT is a fine solution for someone running three Macs and a Linux box behind a cable modem. It does not, however, scale, and This Costs Big Time. If you expect to run a large enterprise on NAT, be prepared to do the moral equivalent of ripping up \$100 bills and flushing them down the toilet, hour after hour.

**Senie:** Fine. Let's focus on your corporate use of NAT issue, since that's where you have problems. I, for one, would like to hear more about how these companies are using NAT, and perhaps we can from there work toward some solutions.

I'd suggested the other day allocation of a block of addresses, to be handed out by a registry in small chunks, specifically for use in interconnects between companies. I had

heard lots of rumblings in the past that this is where NAT in large companies was causing trouble. From the lack of any response, perhaps the problem is otherwise. Since you're having such troubles, and since many of us apparently don't understand the situations where you're getting into trouble, please enlighten us.

I've generally been of the opinion that NAT is a very workable solution for the small office and home network, and questionable for larger networks. Sounds like you're saying the same.

**Metzger:** I don't agree that NAT scales, ever. Its a great kludge if you are using it for a very limited number of things. As I said, letting two people surf the web behind one cable modem with one IP address works on a NAT. Does NAT scale to true home networks, where every light bulb is SNMP manageable? No. Does it scale even to a modest home network? Only if you think the sole application is letting you surf the web on your toaster, which I doubt is the future.

**Jeffery Altman** responding to Senie's remarks about NAT: The New York City Board of Education is using NATs as a security measure to keep their 1000+ schools off of the public network. Teachers are reporting that the networks are unusable because of them. Many of the educational benefits that the schools want to gain from being connected to the internet are unaccessible because of the limitations NATs place on the types of connections that may be made (and accepted.)

The NYC BOE does not have the money or staff to figure out how to properly configure and maintain these devices. But they were put in place most likely because they were presented to the high level admins as an easy way of securing the network. The teachers (as consumers of the technology) have chosen not to use the Internet in the classroom experience because they can't get the same access from the classroom that they are able to receive from AOL at home.

**Metzger:** Actually, to a large extent, the "internet" as "transparent end to end catanet" \*is\* dead. It has been dead ever since the average company was forced to use private addressing. Maintaining the pseudo-internet we've got left is costing us a lot of money and manpower, and neither is an unlimited resource. It will make people very happy if we can give them an internet back again.

**Senie:** The counter argument is that for the Home Networking case, which is a HUGE market, NAT is indeed cheap and easy to use. ... NAT can be used for a variety of things. Perhaps we can agree that it's a good hammer when the nail is a home network,

and concentrate on what to do about the large corporation issue.

**Chiappa:** This relates to a thought I've been having over the last couple of days, which is that I recently read that the Internet usage numbers in many large cities in the US (sorry, no idea about the rest of the world, and in any case this point relates to ARIN only) is now at or over 50% of citizens - i.e. in the US, the logistic growth curve for that group (which ought to be the largest possible market segment) has started to tip over.

So my question is: I've been hearing that ARIN is the stingiest of the registries when it comes to handing out IPv4 addresses - is there any valid reason for this extreme parsimoniousness (particularly when the plan is to move over to IPv6, so there ought to be no reason for extreme hoarding of IPv4 addresses)?

Perhaps a little loosening of the address allocation tap at ARIN, when it comes to allocating addresses for non-home use, could make life substantially easier for the segment where some people are finding NAT making their life difficult?

**Kim Hubbard [ARIN]:** I wouldn't say that ARIN is stingy in how much address space it issues, we allocate whatever an organization (any organization) can justify but they must be able to justify at least a /20 (by whichever allocation/assignment policy applies to them) or they will be referred to their upstream ISP. This isn't done as much for conservation of address space as it is for aggregation of routing table space. We recently lowered our minimum allocation from a /19 to a /20 to allow more organizations to come directly to ARIN for address space and are monitoring the effect on the routing tables to see if we can continue lowering it in the future.

**Keith Moore** to Don Senie: NAT is a good hammer for a home network if and only if the only purpose of a home network is to allow multiple web clients at home to talk to servers in the outside world.

If you want to use a home network to be able to access your devices at home \*from\* the outside world - e.g. IP telephony, IP fax, instant messaging to your home, IP printing to your home printer from elsewhere, setting your VCR, setting your thermostat so that the house will be warm when you get there, checking the house temperature to see if the air conditioner has died again, taking a peek at the kid you've left home with the babysitter (or by himself) to see that he's okay, investigating the alert you got from your intrusion detection system, personal web server for home or home office - NATs start to look like a pretty poor hammer even for home use. (unless, of course, you think

the purpose of hammers is to break things)

On the other hand, if you combine NAT with 6to4 for home networks, the picture starts to look a bit better. Think of 6to4 as the generic ALG that rids you of the need to have separate ALGs for most of the applications that NAT happens to break.

**Tripp Lilly:** Mine is not a stand in favor of NATs, let me get that out first :-). However, the arguments against NATs in the home all center around end-to-end connectivity to various devices in the home (light bulbs, toasters, VCRs, thermostats, etc).

Is this really the "right" model for that sort of interaction? Personally, my home network (in which every light bulb \*will\* be on the 'net within the year) is not something I want end-to-end connectivity to. I'm not saying that's the right solution for everyone, but I think it's certainly worth thinking about as we're designing VCR control and LBMP (Light Bulb Management protocol). That is, I think it's important to consider that folks (via their vendors) will want to deploy ALGs at the boundary of the house, NAT or not. I know I will be, even after the internal v6 infrastructure meets up with the rest of the world in the far flung future.

I don't think NATs are architecturally "correct", but I think they're teaching us an important lesson about the (initially valid) assumptions about end to end connectivity. Even after we eradicate NATs through wholesale migration to v6 (optimist hat on), the paranoid will still deploy ALGs on their firewalls to mediate access to those globally routable lightbulb and security camera addresses. After all, I wouldn't want the world getting illicit shots of me in my underwear in the evenings. Well, perhaps it's the world that wouldn't want to be getting those shots, but you get my point...

**Senie:** Sounds to me like at best I'd trade a NAT box with firewalling for a serious firewall. I have ZERO interest in allowing the kinds of things you describe to occur from outside. While you may not mind someone hacking into the microphone on your PC and using it as a bug I am a little less trusting.

**Moore:** Obviously you have to have some security measures in place before you open up such things to the outside world. but that's an argument for better authentication technology, not for NAT. without the NAT in place I could use IPSEC to authenticate myself and punch a hole through my home firewall; with NAT in place that's just not possible.

Referring to Lilly's remark: 'Personally, my home network (in which every light bulb \*will\* be on the 'net within the year) is not

something I want end-to-end connectivity to”, **Moore** asks: “why not? [It] seems like if you want your light bulbs to be independently addressable or pollable (can’t wait for the SNMP lightbulb MIB!) you want the ability to talk to them directly. OTOH, if for the specific case of light bulbs you want some sort of “light management system”, then maybe you want to talk to that light management system rather than to the individual light bulbs.”

The point is, even though you might want some local resource managers to mediate the use of light bulbs, whatever, that doesn’t mean that you want to block all outside connectivity to every device in your home. If you have to have an ALG for everything you want to control from outside that is going to impose a serious barrier to the kinds of controllable devices you can have in your home - because you won’t be able to control it unless your NAT supports the right ALG for each device you want to use.

And it’s downright silly to have a wireless PDA (say a palm vii or a pdq phone) and not be able to use it to talk to your devices within the home just because your PDA and your home are on opposite sides of a NAT.

Then referring to Lilly’s statement “I think it makes sense to consider a boundary (firewall+ALG) that defines a “trusted zone” within the house, establishes ACLs for a given “connection”, be it a tunnel or otherwise, defined by an authentication event, and mediates the activity over that connection as long as it’s active,” **Moore** responded “you’re confusing trust boundaries with network topology. Trust boundaries don’t follow network topology even today, or you have to do a fair amount of work to make them do so. they’re even less likely to follow network topology in the future when a significant number of the devices we want to talk to are running wireless IP.”

“And just because I have multiple devices in my home doesn’t mean that I trust my (roommate, spouse, kid, babysitter, houseguest, burglar, landlord, friendly neighborhood cop) to have net access to everything in my home merely by having physical presence there. Nor do I want to have to run separate protocols to access devices on my home network than for the same kinds of devices located in other environments.”

**Metzger** in answer to Lilly: I don’t want to invent fifteen thousand different protocols to handle things. IP already does what I need most of the time.

**Lilly**: Perhaps I wasn’t clear... IP (v4 or v6 or what have you) is a fine way of determining the end points of the communication. But at higher levels (MEGACO, SIP, LBMP, etc.), I believe it makes sense to allow in the

protocol design that people might want to consolidate functionality in an ALG (more below)

**Metzger**: I’m not sure that’s the right model, actually. IP addresses are too easy to forge. The right way to stop people from doing that sort of thing is to deploy end to end security protocols that strongly authenticate both ends.

## Uniformity versus Diversity

**Lilly**: Realistically, in the home environment (which is quite specifically the domain I’m constraining these statements to, even though they might have broader applicability), it’s unreasonable to expect that every light bulb (light fixture) is going to carry the silicon to handle authentication (and/or encryption).

I think it makes sense to consider a boundary (firewall+ALG) that defines a “trusted zone” within the house, establishes ACLs for a given “connection”, be it a tunnel or otherwise, defined by an authentication event, and mediates the activity over that connection as long as it’s active. Treating each and every action into that trusted zone as a separate request, carrying separate overhead for connection setup and teardown (over the WAN), and separate overhead for authentication and encryption puts us in the same boat as HTTP/1.0.

I’m not saying we should consider anything other than IP to establish the desired endpoints of the given transaction. I’m not saying we should try to hide topology and addressing behind a NAT. I’m saying that even \*with\* a connection that’s end-to-end for the purposes of designating participants, we might want to consider that someone in the middle will be mediating the conversation, acting on behalf of one or both participants.

An example to wit: I want to be able to plug my Extend-A-Home 2000 (tm) intelligent brick into the Ethernet jack in my hotel room, then unpack all the rest of my goodies (portable printer, portable scanner, wireless IP phone, Palm Connected Organizer(tm), MP3 player, etc.) and have them “just work”. Now, I realize that all of this can be accomplished through a combination of DHCP, DDNS, and IP Mobility. But that requires an awful lot of complexity in each device, when that complexity \*could\* be hidden inside the Extend-A-Home 2000 (tm). I plug it in and \*voila\*, my hotel room is an extension of my home. All of my permissions into my home remain intact (with only an authentication exchange between the Extend-A-Home 2000 and the Home-Weiller 2000 Border Establishment Unit(tm)).

You also have to consider that just because IP is the “right” answer doesn’t mean it’s what will end up in the stacks of all of these micro devices (especially light bulbs). There will be gateways and proxies for LON and CANbus and X-10 and so forth for a while to come, possibly forever. All I’m saying is that taking ALGs into account for reasons \*other\* than NAT doesn’t seem like such a bad idea as we’re doing new work.

**Metzger**: Anything mankind can lock, mankind can unlock. You will never get rid of firewalls. At least not in our lifetimes.

On December 8 **Moore** responded: actually, I’m recently forming a radical opinion that firewalls need to be first-class components of the internet architecture. Only: (a) they should be thought of as “access control checkpoints” rather than as held responsible for authentication (just because you can get through a firewall doesn’t mean you’re authenticated for all services beyond that firewall), (b) we need authentication mechanisms that allow our packets to traverse multiple firewalls (including both ingoing and outgoing firewalls) and still serve as authentication for services at their destination. that is, we need to be able to attach (perhaps multiple) credentials to packets, that stay with those packets end-to-end rather than having to do tunneling. Those credentials (sadly) may need to be based on both user identity and current network location. It should follow that (c) IP addresses have nothing to do with authentication in such a world - there will be too many cases where trust boundaries and IP topology don’t coincide, and trying to do VPN-like things for all of the different things you want to authenticate to from the same host will be too hairy.

**Lilly**: I agree with this... My earlier point about ALG’s wasn’t intended to be in support of NAT (I specifically disclaimed that, in fact), nor was it intended to suggest that trust boundaries and IP topology coincide (even though my examples were of situations where they did). I guess I was being too specific about a perhaps more general problem, which is that of allowing border devices to act as proxies, agents, interlopers, etc., for interior devices. Whether “border” and “interior” are defined in IP topology terms for a given installation or in authenticated identity and trust relationship terms is, I think, immaterial. The point is not to \*assume\* that the actual endpoint (as specified by the IP/port tuple) is going to be the one engaging in all phases of the interaction on its own behalf, and to consider how that non-assumption affects the protocol designs.

That is, NAT has taught us a lesson. NAT is bad, but the lesson is good. Don’t throw the lesson out with the NAT. Furthermore, I’m

not asking anyone to solve the problem of how you maintain those proxies or ALGs or whatever they are. I'm simply asking folks not to build a world in which they cannot work without significant "hackish" after-engineering.

**Ian King:** If you want to be able to control individual lightbulbs in your house, how about an IP <-> X10 gateway? X10 (as an example, not because I have any particular attachment to it) is a useful protocol for controlling devices (usually with limited "intelligence") within a home; that's what it was designed for. Do you use SMTP to program your router? Let's use appropriate protocols at appropriate times and places.

**Metzger:** Insufficient. I want IP. There is no point in implementing anything less powerful than IP — you only end up maintaining a second, inferior stack. IP does what you want already, and better.

**King:** PDA? Why can't I talk to my home machine? I do it all the time (not with a PDA, but with other devices), even though the machine is inside a NAT "boundary". Yes, it took a little configuration magic, but nothing as complex as e.g. RSIP. Why is this about "controlling the household from the outside"? I thought this thread was about "big companies controlling large private address spaces"? Or end-to-end UDP for the latest version of Quake? Or SNMP to flush my toilet from McMurdo Sound? This thread has wandered over a lot of territory.

NAT IS A HACK. Is it a useful hack? In some circumstances, yes. I use it, without a lot of attention on it; my wife uses this computer to browse the Web and read email, and when I say NAT she reaches for the bug spray. :-)

Does NAT work at the ISP level? Depends on what the ISP is selling. I had a lot of problems with ISPs who wanted to sell me a "black box" of Internet connectivity — I'm a geek, I know what I want, and some of them didn't want to sell it to me (in large part because their sales reps didn't understand what I was talking about). Does an AOL user care about his place in the address space hierarchy? Likely not, nor does he "purchase" a right to care. Do I care? Yes, and I pay for it, and I get it.

Does NAT work for corporations? Depends on what they are hoping to buy with it. A NAT client at Foo Corp can attempt to access resources across the net; if Bar Industries is also using NAT, Bar's NAT must be configured to direct the incoming requests. But that can work just fine. Multilevel NAT (i.e. within an organization) is as prone to problems as multilevel marketing; Just Say No. Corporations can use firewalling, internal network addressing, and proxies.

Microsoft does this, and I rarely experience a situation in which I cannot do the "end to end"-ish thing I want to do. (You can't do it from outside, but that's why it's a firewall.) NAT IS A HACK. Why is there so much effort going in to somehow either "legitimizing" it, or demonizing it?

**Metzger:** Because there is a fight brewing about IPv6 and whether NAT is a sufficient alternative to IPv6.

**King:** As I've said before, I use it because my ISP is greedy and wants a lot of money for more than one IP address; I think they assume I'm doing something "commercial". (I also pay a premium for my DSL connection because it's not the base "consumer" speed, and USWest assumes it's "commercial".) As a side-effect, it creates a level of security; my "inside" machines are not directly on the Internet, and it makes it harder for them to be compromised. (Not impossible, but harder, and there's nothing there that makes it worth it.) That's why I use NAT.

Is NAT right for you? NAT IS A HACK. Does it serve your purposes? If so, cool, go for it. If you are a vendor to others, does it serve your customers' purposes? If not, and you are selling them something you can't and don't provide, then you are a crook and legal process should be invoked to deal with you.

Does SMTP give you "end to end connectivity" to each email user you address? No, they can be on completely disjunct machines, with incompatible (or no) network capabilities. (When you send me email, you are NOT sending it to the machine where I read it, nor could you get here from there.) There are certainly some protocols that fail without "real" end-to-end connectivity. There are many that do not.

It is an invalid assumption that, as a class, devices won't be able to communicate with your home devices because of a NAT — in some cases it's true, in others, false, in others, it requires a little more hacking. NAT IS A HACK. Maybe a particular circumstance requires a little more hacking, maybe it requires an ALG, or maybe it requires a redesign of its protocol to allow for NAT. Question: is what you get, worth the effort? Long term?

NAT IS A HACK. Let's step back and focus on ways to fix the problems that led people to think of NAT in the first place, rather than trying to engineer NAT as a long-term solution to those problems. Perhaps NAT will remain as a solution for a certain, smaller class of problems — cool. If NAT isn't solving your problems, DO SOMETHING ELSE. But building a world on NAT is building a world on a HACK.

**Metzger:** I doubt any average homeowner could effectively run a firewall. It is necessary that the devices be secure ab initio, and only communicate to properly authenticated and authorized sources.

**Ed Gerck:** And yet, there is a trend towards "personal firewalls". Linux includes a firewall out of the box (with the ipfwadm and ipchains components). Several products are on the market for Windows — see <http://grc.com/su-firewalls.htm>. One product is very user-friendly, it seems to me that any homeowner could use it.

So, perhaps the same company could also make a NAT that any homeowner could use? Because if the problem of NATs is easy of use, and this is the key being banged here (the NY School Board example, etc.) then it is a problem of design. However, if the problem is concept, in which way are NATs different from gateways, conceptually speaking? And, gateways are useful, no? Further, it seems to me that if NATs are to be blamed for the demise of IPv6, or its ad eternum delay, then maybe this is what the market wants — a multiple-protocol Internet, where tools for IPv4/IPv6 interoperation will be needed ... and valued. A commercial opportunity, clearly.

**Chiappa:** This relates to an approach that seems more fruitful, to me - let's try and figure out things that sidestep this incredibly divisive, upsetting and fundamentally unproductive argument, and try and find useful things we can do to make things work better.

**Gerck:** Which can, undoubtedly, be put in a sound theoretical framework for NATs, in network topology. NATs do not have to be a hack. They seem to have been discovered before being modeled, that is all.

**Chiappa:** Well, the fundamental architectural premise of NAT's \*as we know them today\* - that there are no globally unique names at the internetwork level - is one which is inherently problematic (long architectural rant explaining why omitted). So I don't think that the classic NAT model is a good idea, long-term. However, I think it's a bit of a logical fault to think that the only options are i) IPv6 and ii) NAT's. That's like saying that if my dog's not in my office, he must be on Mars - there are other alternatives! This is especially true in the long run; we may be stuck with NAT in the very short term, but in the longer term we can explore other alternatives.

**Gerck:** So, much as I side with Perry's defense of IPV6 though, I cannot side with a downplay of NATs in order to leave more room for IPV6. Indeed, NATs can help IPV6 interoperate... so, it is by definition, useful. And firewalls are IMO much more home-

owner-friendly than “ab initio security”. So, we need to be careful otherwise the baby goes with the baby water ;-)

**Editor:** Keith Moore however would have nothing to do with Gerck’s propositions. His assessment was that the upper hand in designing a software environment must always go to the engineer as the only one with knowledge adequate to the subject at hand.

**Moore:** NAT’s problem is not ease of use. NAT’s problem is that they break things in subtle ways. Many users can install a NAT, but fixing the problems caused by NATs is beyond the ability of all but the most sophisticated users. (and those who do have the ability would far rather their time not be wasted on such pursuits)

**Gerck:** And, gateways are useful, no?

**Moore:** NATs, backhoes, dynamite, carbon tetrachloride. All of these are useful, in limited situations, by professional experts who know the risks of using them and take adequate precautions to minimize the danger associated with their use. That doesn’t mean you should try using them at home.

**Gerck:** Further, it seems to me that if NATs are to be blamed for the demise of IPv6, or its ad eternum delay, then maybe this is what the market wants?

**Moore:** perhaps. but we should not confuse the market with intelligence, or “what the market wants” with sound design. there is sometimes a rightness to “what the market wants” (meaning that the market is sometimes wiser than widely publicized experts) but the market is not an infallible source of wisdom. and the market cannot choose wisely if engineers and vendors don’t provide it with good options.

If you do a cost-benefit analysis for NATs vs a large flat address space you will almost certainly find that NATs have a favorable short-term benefit/cost ratio (for some cases) and a very unfavorable long-term benefit/cost ratio. This might be fine if NATs are treated as a short term hack or a method of transition to IPv6. But if the market over invests in NATs in the short-term there is some possibility that you cannot reap the long-term benefits of IPv6. The market is not necessarily endowed with foresight (indeed, our economic system seems to artificially and unwisely favor short-term gains), and hill-climbing strategies often do not yield good results.

## Issues Affecting ISPs in the Allocation of IPv6

On December 9 to a question from Sean Doran about IPv6 allocation **Christian**

**Huitema** responded: These entries are supposed to be assigned to “top level” providers — the equivalent of v4’s default-free ISPs. There is indeed a concern that 8K may not be enough in the long run, which has lead to place the next 8 bits in reserve. Each of these “top level” prefixes thus contains 24 bits of address space for organization or delegation by the provider.

**Sean Doran:** So, forgive me for asking a stupid question, but RFC 2374 [Editor: An IPv6 Aggregatable Global Unicast Address Format] is full of oddities, especially in s3.2. Are you of the belief that as a matter of policy, everyone but “top level” providers will have addresses from a “top level” provider, with no exceptions?

**Huitema:** Let’s put it this way: the registries are instructed that only top level providers should get one of these addresses. Everyone who does not qualify supposedly get a delegation from a TLA, or several delegations in the case of multi-homed networks. Who qualifies is indeed an interesting matter of debate for the policy councils of the registries.

**Doran:** Do you also believe that for inter-TLA routing information-exchange purposes, with respect to the destination address, ONLY the 13 (to 21) TLA (+ RES) bits you mention should ever be considered by a router in the core of the global network, except where two directly-peering TLAs agree to exchange some NLA information?

**Huitema:** Yes, absolutely. There should not be a requirement that inter-TLA routing carries anything else, except for bilateral agreements on a voluntary basis. It should be entirely within the specs for a backbone provider to ignore all announcements that are more specific than the TLA bits.

**Carpenter:** This is the plan. We start out with no holes in BGP and we try to keep it that way.

**Doran:** Ah, BGP. Ok, so it seems like there is a 1-1 mapping of TLAs to AS numbers — after all, why would an AS encompass more than one TLA prefix?

What happens if and when the 8k limit on TLAs is relaxed, because the number of organizations qualifying for (however this might happen) a TLA assignment exceeds that number, given that the relaxation might well take us beyond 16 bits of TLA (I take this as implied by Christian’s earlier comment).

The second half of the question above: why would one split a TLA across multiple ASes in the default-free zone? For NLA<>TLA, do you propose to have (possibly floating) static routes be the means of handling NLRI?

That is, we configure default on each operational connection to the TLA that the NLA connects to (and deconfigure at each interface when it is not operating), while we configure static routes to the NLA addresses at the border of the TLA. Yes? I found no FM to R on this subject.

Will any effort be made to keep the NLA IDs both globally unique and portable? Or, because TLA+NLA may share the same NLA bit pattern with another 2nd-level network TLA’+NLA, will 2nd-level networks who multihomed to multiple TLAs do one of: (a) be “promoted” and become a TLA themselves (b) carry multiple NLAs, assigned by each TLA they interconnect with ? If (b) do we make the tunnel hack work, or do we make the NAT hack work, or both, or do we just ignore what to do about losing connectivity to one or more of the TLAs to which a 2nd-level network is connected?

On the other hand, if the NLA bit patterns *are* kept globally unique and portable, how will reachability via one or the other of multiple TLAs be announced to the global core network?

**Bill Sommerfeld:** One thing worth noting is that anyone out there who has a \*single\* ipv4 address has a /48 in ipv6 space already, thanks to the 6to4 mechanism.. It seems obvious to me that the only way routing can scale with addresses this large is with very aggressive aggregation.

**Sean Doran:** It would work better still with abstraction, hence my (maybe stupid) questions to Brian and Christian.

**Sommerfeld:** The only way multihoming can work when aggressive aggregation is in place is if hosts end up with multiple addresses (one from each prefix) and know how to use them intelligently... and, if we can manage to get that to work well, it has the advantage that end sites — even very small ones — will be able to multihomed merely by buying service from multiple providers. It does make life interesting for mid-level providers which want to multi-home, though.

**Doran:** Even trickier: how to get non-local hosts to use them intelligently. The DNS as a mechanism for policy-based routing (in the sense of remotely influencing other people’s path selection) is not totally without appeal. But then again, I’m a loon, and like NAT.

**Huitema:** This is definitely a research issue. I think however that there are at least three possible solutions, and so I believe that this is not a very difficult research issue.

The first solution is indeed what we do everyday: get several addresses from the DNS, pick one more or less at random, try it, and

if it fails try the next one. The second solution is when your DNS resolver has acquired some knowledge of the Internet, and can sort some. Strict provider addressing actually makes this kind of knowledge acquisition slightly easier, as the knowledge table is essentially similar to a routing table, and thus subject to the same kind of aggregation. There are indeed many ways to acquire this knowledge, from looking at the routing tables to getting feedback from the hosts, and this is where research becomes interesting. The third solution requires TCP implementation that have the "zero context" hacks necessary to protect against SYN flooding attacks. In that case, the caller simply sends parallel SYN messages to all possible addresses, continue with the first response, and forgets anything else. Basically, you trade routing complexity for increased traffic. A related problem is the support of readdressing, i.e. changing the IP addresses without losing the TCP connection, but there are already solutions in the v6 spec.

**Editor:** Earlier **Huitema** had said 'Let's put it this way: the registries are instructed that only top level providers should get one of these addresses. Everyone who does not qualify supposedly get a delegation from a TLA, or several delegations in the case of multi-homed networks.'

**Moore:** Of course, this requires that sending hosts or applications make intelligent decisions about which destination address to use (and which source address to use with a particular destination address), usually in the absence of any information which might inform the decision.

It's not at all clear that this can work well enough to be a general purpose multihoming mechanism, at least not without adding a fair amount of extra infrastructure and complexity - i.e. a mechanism which hosts or applications can use to query the network to determine relative proximity of several different addresses. If it does turn out to work it will probably be because all of the available prefixes for both the source and destination host are so reliable and have so much available bandwidth that most of the time that it doesn't matter which of the available addresses you use. (It's tempting to say that multihoming will work quite well for those cases where you don't need multihoming... but that is a bit of an exaggeration). To be fair, "traditional" multihoming doesn't scale well enough to use that approach either.

On December 10, **Jessica Yu:** There is also a potential scaling issue of using multiple addresses as general purpose multi homing mechanism. This is because if this is the case, most of the Internet hosts will end up with multiple addresses.

Based on the current assignment strategy,

only top-tier ISPs (networks) will be assigned with TLAs. Tier-2 ISPs will get allocation address space from the TLAs of their connected tier-1 ISPs. If a tier-2 ISP is multi homed, the ISP will be assigned with two IP addresses from two different TLAs. All the downstream customers of this tier-2 ISP will get multiple addresses for each host. Since most of the tier-2 ISP (if not all) will be multi homed to tier-1 ISPs for redundancy purpose, most of the Internet hosts will end up with multiple IP addresses. Only tier-1 ISPs' direct connected single-homed customers will not have multiple addresses for their hosts.

It's possible that some multi homed sites will have to assign 4 or even more ip addresses per host, depend on what kind of ISPs they multi homing with. E.g. a site that happen to multi home to two tier-2 ISPs, each multi homed with two different tier-1 ISPs, each host in this multi homed site will have 4 IP addresses in order to get full benefit of redundancy. One can imagine even more complicated case. So I think we need to explore different ways of doing multi homing.

**Sommerfeld** (to jessica's statement of most hosts ending up with multiple addresses): I don't see why this is inherently a problem.

**Jessica Yu:** This is paradigm shift in the Internet from majority of hosts with single IP address to the majority of the hosts with multiple IP addresses. Many existing support mechanisms such as routing (see Keith's message), DNS name look up, traffic engineering network management, etc. may not be adequate. It may also break the things that we have not even thought of. And do not forget about operational complexity issues. Are we really ready for such a major shift? So I would not say so quickly that it's not a problem.

**Senie:** I can imagine that a company which is multihomed to two or more ISPs, each of whom is tier-2 or below, may well have a dozen addresses to deal with, per host. Now somebody remind me of why we wanted all this extra address space? It was so we could give every machine a dozen addresses? Exactly how the end nodes are to know which of these addresses to use, especially when the decision point in the topology is several layers above (at the tier2 to tier1 provider attachment point) is going to be interesting.

Token-ring source routing allowed relatively dumb bridges to be used. This protected the "core" of the network from having to have extra processing horsepower. Of course we soon saw Ethernet bridges and switches which could handle all of the needed decisions without the involvement of the end stations. We started down the path of MPLS to protect the core routers which were going to melt under the traffic load. Then ven-

dors showed us hardware which could route packets at wire speed. Now we're busy pushing technology to ensure routing table sizes won't increase because the present generation of hardware can't handle larger routing tables. Are we being a bit short-sighted?

**Chiappa:** No. The issue with routing table size is not the memory required to store it, but the stabilization time when a topology change happens. That is a factor of a number of things, including the speed of light, which isn't going to get faster any time soon... :-)

Stabilization time is larger in a larger network. Full stop. A larger network will \*also\* have more topology changes per unit time. As a net gets larger, you are caught between two closing jaws: increasing stabilization time, and decreasing MTB topology changes. There are solutions, but throwing more processing power in the routers at the problem isn't it. There \*are\* some ways to deal with these problems (e.g. by bounding the number of nodes that updates have to propagate through, by careful design of topology and abstraction hierarchy, and limiting inter-level linkage) but they demand a more engineered (on the large, inter-provider scale) net than we have now.

**Sommerfeld:** Given how hard it is to get an ISP do to anything special for you these days, I really can't see a routing-system-based multihoming actually scale down to, say, individual SOHO networks being multihomed, while multiple-address-based multihoming doesn't require anything special out of the ISP's..

**Moore:** no, I can't see how a routing-system-based multihoming will scale down to SOHO networks either. but neither do I think that DNS based multihoming will work well, at least, not without developing a lot more infrastructure than is even on the drawing board at present. and making this work well in practice is at least a few years away.

On December 11 **Doran:** Why not? It should be an explicit goal of the next cut at an IPng, and shouldn't just stop at SOHO -- it should go right to individual homes, \_at least\_. I have multiple sets of wires coming into my flat, and live in a regulatory environment where each of those wires (and some new ones) is allowed to carry data/POTS/ISDN/cabletv/you-name-it.

And then there's Sweden, which is several steps ahead. Check <http://www.bredbandsbolaget.se/eng/node103.asp?ID=16> Yes, they really are serious about more than a quarter of a million homes which will have the moment-to-moment chance to use any or all of: FTTC/FTTB/FTTH from these folks, cable modems, or the various offerings from the various telcos (xDSL, dialup) for Internet connectivity.

The days when only a handful of middle-aged Swedish women have broadband connections like this: <http://www.stupi.se/Bilder/7296-3251-0759/g.html> into their apartments are numbered. (BTW, if you look closely, you'll see her apartment's network was multihomed already.) The assumption that only larger organizations will want multihoming and at-will-provider-changing is a bad one.

On December 14 **Jessica Yu** to Brian Carpenter: Looks like we have a terminology issue. Notice I did not say routing system but ROUTING will have problem. Because the choice of a multi-addressed host to use one of its IP address to include in packet header implies routing decision, the host, in effect, does some routing decision making. How a host intelligently choose the IP address as the source address to send packets or DNS intelligently choosing an IP address for a query are complex issues which we do not have an answer for as you indicated in your message.

**Huitema:** I believe you overstate the problem. To begin with, a domain manager is never obliged to advertise all possible routing prefixes for all its hosts. In principle, it will only advertise those network connections that are deemed "good enough." So, at a 10,000 feet level, picking an address at random in the list, without any information whatsoever, gives you a connection that may

not be optimal, but is probably reasonable. Which means that the proverbial "five lines client", which presumably will transfer 5 packets, will get adequate service. You will tell me that this may not be optimal, but then, explain me exactly how the current BGP fabric guarantees that routes are picked optimally? In fact, in their communication at the last ACM SIGCOMM conference, "On Estimating End-to-End Network Path Properties", Mark Allman and Vern Paxson observe that it is quite frequent for interdomain routing to converge on a less than optimal route today.

The question indeed arises when one of the addresses in the list goes through a provider that is experiencing heavy congestion. Let's first remark that, if clients were to pick addresses at random from a list, we would obtain automatically some level of traffic spreading, which would generally tend to ease congestion — but we all agree that congestion will not be eliminated. If the congestion is already present at the beginning of the connection, the answer is simple. The first SYN packet gets lost, and the client simply picks another address in the list and tries again. Again, this may not be optimal, but it will provide an adequate solution for even the most brain dead of clients. The smart clients will no doubt find how to resolve the problem in a smarter way.

So let's focus on the remaining part of the problem. What happens if a client picks an address at the beginning of a connection, only to find out that the transit networks are becoming congested. There are solutions, which require minimal adaptation to TCP. An example is discussed at <http://www.chem.ucla.edu/~beichuan/etcpc/>; another solution is developed as part of MobileIP. Indeed, if there is no modification to TCP, there is no good solution. But then, this is not very different from the current state of the art. A TCP connection can only survive the loss of a transit network if BGP manages to detect and correct the loss faster than it takes for the TCP timers to break. If I believe the recent reports to NANOG, this is far from guaranteed today.

So, yes, we have a problem. We need to somehow adapt the TCP stack to survive losses of transit networks. But we should not overstate that problem. It only affects long connections, it only makes a difference if a connection to a transit provider breaks, and if the routing tables could have been repaired by BGP. In any cases, there are simple modification to TCP, for which we already have some experience, that could handle the problem. In the long run, once these modifications are in place, we are better off than in the current situation, because we will rely less on the speed at which interdomain routing converges.

## Farber Moves to FCC as Chief Technologist

An interesting Clinton Administration gambit is underway. On Monday January 3, 2000 the FCC announced. David J. Farber, currently the Alfred Filtler Moore Professor of Telecommunication Systems at the University of Pennsylvania, has been named Chief Technologist for the Federal Communications Commission (FCC). On the same day Farber explained to his IP list: "I will be at the FCC as an IPA -- a US government mechanism which allows me to serve the country while remaining a Penn Faculty member. I look forward to returning to Penn in a year or so with a great deal of experience and knowledge regarding the policy and economic issues involved in deploying the new communications technologies."

Given our remarks about Farber in the January *COOK Report*, readers will not be surprised to hear that we were less than thrilled. On the Cyber telecom Law list we wrote: my cynicism asks whether in this day and age a consummate insider can be trusted also to be a public interest advocate. You seem to aspire to be both. And I suppose, to the extent you can bring it off, more power to you... I worry sometimes when I have criticized you that I may be too harsh... I mean maybe you do stay up nights worrying about whether ICANN is out of control. Maybe you really aren't helping them out by being a

sometime critic? How do I know for certain and the answer is of course I don't.

Yet you *have been* one of the most consummate insiders in this industry. Witness your role in the July 30 IBM - NSI summit meeting that was not known to have taken place until last month. Witness the one of the "25-most-powerful-people-in-networking" *Network World* accolade of a year ago. Witness some recent evaluations of you as someone who has developed extraordinary skills in becoming a confidant of major IT CEOs who has never betrayed a confidence. Witness an assertion that you are so good at this that you carry the corporate secrets of companies like IBM, HP, Intel, SGI, Microsoft and Sony around in your head and and that you have been able to compartmentalize them to such an extraordinary extent that none of the companies feels that your knowledge has been used to their disadvantage. A remarkable achievement. And one that may go far in explaining why, except perhaps for the Microsoft trial, you have generally shied away from taking a position sharply opposed to the interests of a major IT company or a move in the industry to set up something like an ICANN.

Unfortunately, from my point of view decisions made behind closed doors are gener-

ally suspect. Given the money and power at stake in the current reshaping of the infrastructure underlying our entire economy, I cannot believe that you or anyone else is saintly enough to be trusted never ever to act in such a way that the interests of a single corporation or group of corporations is attended to before those of the public. Besides it is all too easy to build an ex post facto justification that what's good for IBM is good for the US. Or for IBM substitute any number of other players.

An hour later the phone rang. It was the high ranking FCC official responsible for Farber's appointment. Gordon: I want you to understand that we are bringing him to town PRECISELY because he is so well connected. We got a believable story about the FCC desperate to do the right thing for the internet.. "Farber understands and he can get through to the top and educate the people who need to be educated we heard." Could be and we hope indeed that it really is. Farber indeed can do good from this slot. But he could also, if he so chose, sell the Internet down the river from the very same position. It would be nice to have faith that this is a positive move - one that can be trusted. After all we trusted Ira Magaziner. How foolish and naive to do so. This time will the result be any different?

## Executive Summary

### Wireless Takes off, pp. 1- 6, 10

We interview Ira Brodsky CEO of Datacomm Research and author of books on wireless communications. We survey the 1999 explosion in digital wireless technologies summarizing with particular attention to their impact on the Internet mobile, fixed and wireless LAN technologies. Brodsky points out that because of advances in digital technology wireless broadband access to the Internet has become a reality. This means that virtually everything we do with wireline Internet connections we will also do with wireless.

The interview explains from a technical point of view what is done to achieve high bandwidth using TDMA and CDMA it covers Triton network Systems and LDMS. It shows through a discussion of Phone.com and the WAP protocol how cell phones are becoming web browsers. It describes PCS as well as Sprint's leadership in this cellular technology. It explains how Metricom intends to compete with Sprint PCS.

Looking at the Europeans it describes TDMA, CDMA and GSM applications in Europe. GSM is an enhanced form of TDMA that is popular primarily in Europe. However many people doubt that it would stand up well against a well aimed rollout of CDMA.

Wireless LANs are critical to the hopes of home networks of IP aware appliances. Costs are approaching \$100 a node and speeds are approaching Ethernet. The viability of this market however is likely dependent on the outcome of the IPv6 deployment discussed in the IETF debate article included in this issue. It may also be impacted by how the network continues to scale as broadband moves into the edges of the network.

With broadband wireless as an end user option, one moves into the reality of access to the Internet being available from any where at anytime under almost any condition. Given the cost and time necessary for the installation of fiber-based local loop infrastructure wireless is becoming a more and more viable local loop alternative. On January 5, 2000 Advanced Radio Telecom announced deployment of 100 megabit per second IP network to connect high speed business LANs to backbones across the US. The ART broadband MANs will be deployed by years end in ten cities across the US. They will use Cisco supplied Ethernet routing and switching products and configure its MANs in a self healing ring architecture "capable of providing 200Mbps of total bandwidth on its bi-directional paths."

According to George Gilder, the growth of

these systems will be sustained by the introduction 12 to 24 months from now of chipsets based on Qualcomm's 2.4 megabit per second HDR data transmission technology that is a flavor of TDMA running in unused CDMA channels. According to Brodsky: "HDR is CDMA; there may be some time-sharing going on, but it would be misleading to call it "TDMA." it would be more accurate to say HDR runs on separate channels that can be either in the cellular/PCS spectrum or outside that spectrum; saying "unused channels" suggests it borrows channels from the voice system." Cellular coverage in the US has evolved to the state where basically any customer can roam nationwide by using an analogue phone that is also either CDMA or TDMA compatible. Finally Brodsky is optimistic about the capability of a company like Cisco to sell wireless spread spectrum equipment to ISPs who could use it to by pass the LECs local loop strangle hold.

### 2000 State of the Internet, pp. 7 - 10

In our annual State of the Internet essay we examine the continued triumph of the Internet in the areas of wireless and broadband technology, in electronic commerce and in newly evolving data storage technologies. However we caution analysts not to find a false sense of security by restricting their analyses to just these areas. For two areas not involving technology impact the Internet's future. They are the drive to regulate and control by means of ICANN as evidenced especially by those companies still dependent on the success of their legacy based technology. They also include an argument over architecture that appears as a commitment to make wide spread deployment of IPv6 a reality.

The IPv6 commitment is part of a technical debate over what some perceive as the lost "potential" of end-to-end IP connectivity as NATs and firewalls have come to shield or otherwise protect hosts on corporate intranets and prevent several important protocols like IPsec from penetrating the NAT and firewall barriers. The way these concerns are handled will affect the structure of the Internet. It will be extremely difficult to make progress on the implementation of IPv6 without a centralized top down drive designed to get as many as possible to change. But the amount of attention and effort given over to this drive will impact the final area which is one focusing on a slowly growing concern over the continued scalability of internet architecture and routing as broadband technology is deployed at the edges of the network.

Some find it worrisome that some of the scalability issues such as competing backbone architectures examined in the January 2000 COOK Report are generally not discussed openly. These folk believe that these issues may be more important to the smooth future functionality than the outcome of the IPv6 deployment issue. If too much emphasis is placed on new technologies regardless of their impact on network architecture, network performance is likely

to seriously degrade. If too much emphasis is placed on the struggle to control through code or architecture in addition to law in the way that Lessig points out in his Code and Other Laws of Cyberspace, the ability of engineers to handle the challenge of architectural design issues will be severely impacted.

Consequently the overall success or failure of Internet architecture will be determined by the interaction of these three with each other. We contend that most analysts are aware of the technology issues and are making the mistake of focusing on them to the exclusion of the regulatory control and network architecture and protocol design issues. The result hinders the Internet's ability to respond to the demands placed upon it by runaway growth. Successful analysis now demands an ability to synthesize technology, legal and network design issues.

### IETF Architecture Debate, pp. 11 - 24

During the first half of December, on the general IETF list, there was an outstanding discussion of some critical problems of Internet architecture. Most participants were among the most distinguished engineers in the IETF. The focal point was over the dilemmas posed by the desire of these people to gain a set of perceived benefits from the deployment of IPv6. Brian Carpenter's December 1999 Internet draft <http://www.ietf.org/internet-drafts/draft-carpenter-transparency-05.txt> on Internet transparency provided the foundation for the discussion.

The crux of the perceived problem is that in order to make IPv4 addresses scale during the Internet's take off in the mid 1990s architectural "kludges" such as private IP addresses for intranets hidden behind Network Address Translation (NAT) boxes and firewalls and Classless Inter Domain Routing (CIDR) were instituted. The result has been that huge investment have been made in equipment and architecture that will not be easily changed. Also protocols designed to work in an internet with end to end transparency will not work in a world where to get from the backbone to a receiving device on the edge of the network they have to travel through NAT boxes and or firewalls.

The perception is that the kludges are now very cumbersome and costly for corporations to manage. There is a perception that IPv6 which has several orders of magnitude more addresses than IPv4 will provide the Internet with enough flexibility such that the irritating kludges standing in the way of end to end transparency can be removed. Alas this is really true only if IPv6 can be massively deployed through out the Internet. We are talking deployment at such a level that IPv4 virtually disappears. The problem facing the Internet is that ,short of an unprecedented regulatory decree that commands massive adoption of IPv6 globally, enough deployment

**continued on next page**

Continued from page 25

of IPv6 to ever make a difference is unlikely to happen.

Some strong philosophical issues of design and management are at work here. On the one hand the IPv6 advocates have a top down vision of a uniformly designed and managed Internet. Opposed to their view is the belief that certainly reflects the operational reality of the net - namely that the market place is working with the development of diverse solutions that perform quite satisfactorily.

When Ian King wrote: NAT IS A HACK. Why is there so much effort going in to somehow either "legitimizing" it, or demonizing it? Perry Metzger replied because there is a fight brewing about IPv6 and whether NAT is a sufficient alternative to IPv6. Ed Gerck summarized an opposing point of view. Further, it seems to me that if NATs are to be blamed for the demise of IPv6, or its ad eternum delay, then maybe this is what the market wants - a multiple-protocol Internet, where tools for IPv4/IPv6 interoperation will be needed and valued. A commercial opportunity, clearly.

Part of the fight is over control. Who gets to set the rules by which Internet architecture will run? It could turn out to be unfortunate when others believe that there are serious unresolved problems with routing architectures that the time and talent of the IETF is focused on the IPv6 control issues. We may be certain, however that the IPv6 controversy is extremely important and will not quickly disappear.

### Farber Goes to FCC, p. 24

On January 3, 2000 Dave Farber was appointed Chief Technologist at the FCC. Reaction was generally favorable that the Agency would have an Internet expert in the position. We wish that we felt about Farber's mission, as comfortable as the other experts.

### Forthcoming

We intend to publish our annual anthology tentatively titled The Disruptive Internet: Triumph or Chaos by January 20. The March Cook Report should appear before February 1. It will contain an interview on Gigabit and ten gigabit Ethernet and an interview with Kathy Nichols on diff serv.

## IP Insurgency: Internet Infrastructure & Transformation of Telecomm - \$495 Now available.

### Subscription Rates

Choice of either ascii or Adobe Acrobat (PDF) format 1. Individual; College or University Department; or Library; or Small Corporation - \$250 2. Corporate - (revenues \$10 to 200 million a year) - \$350 3. Large Corporate- Revenues of \$200 million to \$2 billion per year - \$450 4. Very Large Corporate- Revenues of more than \$2 billion per year - \$550

Site License: The right to distribute ascii and PDF via email to all employees of corporation. 5. Small corporate: \$450 6. Corporate: \$650 7. Large Corporate: \$900 8. Very Large Corporate: \$1150 . Site License Distribution via intranet web site \$400 a year additional. See [www.cookreport.com](http://www.cookreport.com) for more detail

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

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