The COOK Report on Internet Protocol
Technology, Economics, and Policy

Building a New Communications System for America at the Grassroots Level

A Vibrant, 21st Century Community Communications Ecology Requires Physical Infrastructure, Human Skills and Social Tools

Editor's Introduction

I spent much effort in November and December helping to gather new information for the Obama transition planners. When I saw Free Press' ill-defined December 17th 44 billion dollar broadband proposal and ITIF's magical "promise" that, for 30 billion, 994,000 jobs would be created, I was feeling glum. The incumbents, it seemed, were all lined up with hands outstretched to "build broadband." Would some kind of sanity prevail?

Finally, it did! On the evening of January 14 Mark Cooper posted a link to an outstanding plan that he and Gene Kimmelman had authored and on the 15th when the overview of the requested spending bill was released the summary text said: "Broadband to Give Every Community Access to the Global Economy · Wireless and Broadband Grants: $6 billion for broadband and wireless services in under-served areas to strengthen the economy and provide business and job opportunities in every section of America with benefits to e-commerce, education, and healthcare. For every dollar invested in broadband the economy sees a ten-fold return on that investment."

This does give some hope that Barack gets it in this critical area. I am devoting the usual interview slot of the March *COOK Report* to the short proposal and attendant discussion. The hope here is that we may be able to avoid the TARP fiasco by not giving billions to the carriers for "broadband" only to see them pocket the money and continue their misguided ways. The hope also is that we may be able to spread some trans sectoral thinking into looking more carefully at a national fiber infrastructure for its use in things like science education at all levels and spreading learning collaborations with payoff in environmental and energy areas.

When I asked Mark for his take he responded: "If making law is like making sausage, then explaining and understanding how it get done is murky. Taking credit is dubious at best. That said:

(1) Opposition to the big spend, gold plated approach

Volume XVII, No. 12
March 2009
ISSN 1071 - 6327

On the Inside

Please read Explanatory Note page 97

was expressed early by CFA and CU to the transitions team.

(2) There were large corporate entities who were opposed to it as well and independently supporting the general approach

(3) Because of the way the financial sector bailout has unfolded, the Congress is ultra sensitive to charges of corporate welfare.

(4) The more money you are spending, the more the Hill gets into the protection of the jurisdictional imperative. There were differences of opinion over who would spend it and how.

(5) When things become controversial in a bill that is supposed to be fast-tracked, there is a tendency to shrink the size and fall back on the noncontroversial. A small amount of money spent on unserved areas is the safe thing to do.

In the hot house atmosphere of the stimulus package the most important thing is to avoid bad policy. That is the outcome, pending Senate action. The next challenge is to actually implement good policy."

Editor: The article by Mark Cooper that follows is found here:

http://www.huffingtonpost.com/mark-cooper/building-a-new-communicat_b_157899.html]

Building a New Communications System for America

"The Washington debate over Internet funding in the economic stimulus package provides a remarkable opportunity to build a 21st century communications system. But if we're serious about making it work, this new communications system must focus on the people and the tools they use to communicate, instead of how fast or large the system is.

How do we build a new communications system of the people, for the people? Look at what the Obama campaign built to communicate with voters. They trained almost 150,000 activists and mobilized 1.5 million volunteers to revolutionize political campaigning in America. Viral organizations turn traditional political organizing on its head. Political parties are focused on finding ways that the members can help the party, developing structures that enable the party leaders to give marching orders to the party faithful. **Viral organizations do the opposite. The viral organization serves the members, giving them the tools to self organize, empowering them to do what they think needs to be done to accomplish the overall goal through local, autonomous action.** The tools are communications devices and networks, hardware and software that allow the volunteers to find and communicate with like-minded people in their communities.

Communications companies are lobbying for billions of dollars in the stimulus package. But what the new administration should be focused on is **how to give people access to the Internet on open local networks managed by cities and counties**.

We envision a community-wide fiber network linking all local government buildings, schools, and libraries. The service would be anchored by local government. Non-mobile communications flow over the fiber network. Mobile communications flow over the fiber network to a WIFI/WIMAX wireless network.

The schools and libraries can also be "hot spots" in a WIFI/WIMAX network that would also be available to the community for broadband communications. As a city, county-based network, the service is provided at cost to consumers and the network is operating in a nondiscriminatory manner.

The stimulus package can be used to create a team -- an "E-Corp" -- to train commu-

nity members in digital communications and digital skills. They can retrain unemployed workers with digital skills to become local tech support. These activities foster the skills for a more competitive work force.

This approach can also address the digital divide that has grown in the past decade. One-third of American households still do not have the Internet at home. Almost half do not have a broadband connection. The biggest problem is the cost of broadband, but there are also skill and attitudinal barriers to adoption. Our community-based initiative addresses all the major obstacles to the adoption of broadband.

All residents of the area served will have access to e-mail, chat and browsing at no charge. Repeaters and hardware are subsidized for low income households. Unserved rural and urban areas receive top priority.

The community network should develop and deploy social networking tools working with members of the community. The implementation of existing social networking tools in the community is based in the schools, local civic organizations and local Chambers of Commerce. Software and training are the activities that need to be funded. Community projects can also produce content and activities that are relevant to and attract the interest of local people.

Funds can flow through four categories of non-profit entities - local governments, cooperatives, non-profit community groups, and public/ private partnerships The public entity can fund public private partnerships and local government. Cooperatives can be funded through the Rural Utility Service. Non-profits can be funded by re-instituting the TOPS program and designating new money for these purposes. If spending money quickly is the objective (not a particularly good one), there is no shortage of civic institutions that could be used to disperse the funds; but the ultimate goal should be to create viable and sustainable communications assets.

The raging debate over how to define broadband for purposes of "special" treatment in the tax code is a dead give away that stimulus spending directed at the big communications corporations can easily turn into corporate welfare. The corporations will use the tax breaks to pay dividends, increase executive pay, or fatten the balance sheet. The way to avoid this trap is to direct funds to local governments and community-based organizations.

This is also the ideal moment to redefine what government can and should do for the people. Providing for the basic means of communications -- paving the streets and building the on-ramps for the information superhighway -- are proper local government functions. The big communications corporations can be hired to dig the trenches the way contractors bid on road and bridge projects, but the people should own the networks and should build the basic communications network that all households need. The private sector can still build its gold plated, hundred megabit network, but it will do so only if people are willing to pay for it. City streets and county roads are open to the least expensive compact car and the most expensive Rolls Royce providing access to basic services for all."

[Mark Cooper is the Director of Research for the Consumer Federation of America. Gene Kimmelman is the Vice President for International Affairs of the Consumers Union.]

Discussion

Mary Beth Henry: Thank you for your excellent post. NATOA (www.natoa.org) has been advocating for local governments to be eligible

for BB funding in the Stimulus bill on Capitol Hill since the week that Obama was elected (actually we have been promoting publically-owned FTTH for many years). We believe, like you, that BB is the 4th utility and as such should be universally available, ultra high-speed and have open architecture. I would be interested in exploring how we can work together to make this vision a reality.

Barron: That's brilliant Mark!

One additional suggestion: the creation of a national backbone [insert joke about growing a spine here] using the NLR. The analogy is often made to the national highway system - here is the opportunity to build the national backbone.`

Cooper: A national backbone would be fine as long as it adheres to the principle that it is affected with the public interest, not a private road. The streets and highways are the most open access networks we have.

Kushnick: I think that Mark's approach is better than most, but unfortunately, the problem I have is - accountability.

You can say I got stuck holding the bag on broadband history in America because I was formerly a consultant and my friends were on the various broadband projects --- and over drinks and non-disclosures it was clear that the companies simply were gaming the regulatory system - which I've been documenting since the formation of New Networks in 1992- thus the name -

I say all this because all proposals to date want to circumvent the incumbent 900 pound gorillas -- I would argue that based on our previous email, we should be getting back the billions taken and reopen the networks. Fios would not only be open but is essentially part of the switched phone networks - as at least the law in NJ would claim. The FCC's rulings on closing the networks never examined the financial contributions or deployments in the states before the closed down rights of opening the last mile, as the Telecom Act's promise turned into mush. We also have the mergers -vertical integration and owning the various components gives the ultimate ability to harm competition - Did anything change since 1984 where we actually cared about anti-trust issues? And on the merger conditions, AT&T was supposed to DSL at 200Kbps or better in all 22 states for $10.00 - They never did any advertising… let's get them to redo their offer and refund the money of anyone who signed up last year and didn't get the discount. _

The problem is --- there are no trustbusters or even those who are going to do anything dramatic pertaining to fixing the core problems. .. Fred and I and some others believe a structural separation is probably the only thing that will fix this mess Very fast open utilities. - which was actually the model promoted by the states in these alternative regulation plans.

Cooper: While you were a consultant for the industry foisting this boondoggle on the public, I was an expert witness for consumers trying to stop it. We agree entirely that the consumer was ripped off. Reminding people that we paid for the network once is a useful exercise, to make sure we don't be ripped off again. While it is laudable that you want to get the money back, fuggedabodit; it ain't gonna happen.

Cecil: But typically the consumer groups are just as unrealistic; they want everything for nothing.

Goldstein: We can't have progress *and* demand a dirt-cheap 1FR rate (single flat rate line) at the same time. We've suffered enough at the tyranny of the low 1FR rate, especially when anti-competitive regulation is justified on those grounds. And

that goes back to Carterfone, if not earlier.

Cecil: The deeper problem is that you were playing within a system where the result was already pre-determined; the framing was wrong. That's also why I'm so deeply suspicious of the muni-bells springing up; they are the exact same business model having all of the same incentives, but might do better because they have the backing of city hall.

Cooper: Accountability, in the current environment means making sure they spend the money for the things they say they will spend it for, but from my point of view, that does me no good. It leaves the incumbents in control of a network that does not meet the real need for communications (networks designed to sell service to consumers, rather than allow citizens and people to speak). Structural separation is a perfectly fine approach, but it does not get you affordable basic service. If you let them build 100 mega bit fiber systems, you will not have enough competition to discipline market power. You have to also regulate the deployment and rates of the transport network. The chances that I will be able to get back to that kind of regulated structure are even lower than the chances I will get the money back.

So this proposal for a community-based, wireless open local network with a robust backbone and middle mile in the public domain to provide as basic service gets the job done and addresses a need (mobile high speed) that I believe has legs.

Cecil: A single fiber optic strand has the potential for more capacity than all of EMR spectrum. And Muni WiFi has been a miserable failure. Why waste the time and effort on this? Moreover, if you are going to build your own backhaul and middle mile, then why stop there? The more likely outcome is that all you'll do is hook those radios up to Bell backhaul, thus killing any pricing advantage at all. Think there's a reason Starbucks went to AT&T? This is frosting over the top of something other than cake. More deeply, until we get the regulatory framing right, every success will eventually backslide down the scarcity slope into existing conceptions of "natural monopoly" - i.e. "it's my dirt" - and we'll be right back where we started. Truth be told, we never left.

Cooper: **You stop at backhaul and middle mile because as I understand it, the cost of the last mile fiber is vastly greater. The prospect of a sufficient number of last mile fiber networks to provide vigorous competition (at least 4, more realistically 6, and in an ideal world 10) is virtually zero. The last mile may not be a "natural monopoly," but it certainly is not workably competitive.** The performance of muni WiFi is a function of the financial models imposed on it by the incumbents. Building community-networks with stimulus funds under a streets and roads model is a very different approach. Mobile computing, local government communications, and basic service are a series of functions that can be met with the community-based network and ensure its usefulness. What is the regulatory framing you have in mind that would accomplish these functions, without building this network?

Coluccio: Mark stated: "The last mile may not be a "natural monopoly," but it certainly is not workably competitive."

Agreed. If the last mile were workably competitive it would have to be, in addition to regulated as such, open to the extent that it would support a COTS-like supply chain. This is clearly not the case. Consider the interesting contrast here.

The incumbents will bring all of their buying power to bear

on manufacturers of WAN element and component manufacturers in order to ensure that they use recognized industry standards that support interoperability among one another, because it suits their larger compatibility needs for aggregation and handoff in what is, by its very nature, a multi-carrier environment. I'd like to declare that the last mile was a fractal representation of the WAN, but for many attributes concerned with interoperability, at least, as things stand today it clearly is not.

In the last mile, the dominant service providers, namely the RBOCs here in the states, and probably their counterparts abroad as well who engage in RFP consortia, have used all of their buying power to ensure that platform manufacturers provide what amounts to "closed" systems, often by virtue of a given model's natural exclusion from mainstream solutions elsewhere in the universe, thus ensuring they meet closed architecture criteria that are designed for maximizing customer lock-in and warding off competitors who would otherwise compete with similar equipment if it were available.

Harrowell: Why the obsession with IEEE802 technologies? What's GSM done wrong other than unwiring the planet? Surely not a tad of Not Invented Here syndrome?

After all, as the then CTO of Motorola Networks said a few years ago, the details of LTE, CDMA 1xEVDO Rev C/UMB (now RIP) and Mobile WiMAX were getting very similar - IP transport, flat architecture, OFDMA and TDD in the air interface, 2x2 MIMO in the antenna, target data rates around 10x10 Mbits.

Goldstein: Just to build on what Mark said, I don't think WiFi (802.11 family) is the answer, but neither is CMRS radio (GSM, LTE, CDMA, etc.). Certainly in the US, CMRS is the exclusive province of the big carriers, the ones who can pay top dollar at spectrum auctions in order to "bank" spectrum to keep out competition. GSM itself is of course an obsolete voice technology; GSMA's 3G (WCDMA) is temporary, and LTE will probably become nearly universal in a few years. But those are all for reserved paired spectrum. They are not available or suitable for community networks.

The mobile carriers (the current FCC has been working to reduce this to a three-carrier market in the US, but right now it's usually a four-carrier oligopoly) are not ones to provide open Internet access. The mobile world, after all, is the main consumer of DPI. They put onerous Ts and Cs on wireless access and the still FCC smiles on them. In contrast, when cable very, very gently enforced its long-standing terms and conditions, they come down like a ton of bricks, even though they had no statutory basis to act. Remember, for the past eight years, the US has had a "government of men", not a "government of laws", and the damage will not be easy to undo.

WiFi itself is an awful technology for this purpose, though. It is indeed a LAN, and should be left for that. Trouble is, there isn't much spectrum for anything else.

Harrowell: In fact it looks like HSPA is getting to those rates already, iteration by iteration, and a lot of products already push the Radio Network Controller down to the Node-B and provide for the breakout of Internet traffic onto the Internet at the backhaul level. This is a sort of implementation of the 3GPP Systems Architecture Evolution proposals and for that matter of WiMAX's network architecture.

And even old GSM catches up; I think the latest EDGE products now break the megabit, you can get lightweight and deployable BTS gear for cheap, and with SIGTRAN and OpenSS7 the rest doesn't have to be ex-

pensive. After all, these radio layers actually work and we know they work; but the world, and especially the United States, is littered with dead public Wi-Fi projects. Ask Earthlink, Google, Philadelphia etc.

The reason is that Wi-Fi is a silly name which causes people to forget that it's really called WLAN, and the LAN bit means what it means everywhere else you meet it. It doesn't mean wide area network, it wasn't designed for it, and it doesn't work well, and the sheer numbers of access points required to provide a decent service kill the economics.

Now, WiMAX is a different kettle of fish, being designed from the start as a MAN technology and now usually deployed in licensed 2.5-3.6GHz spectrum rather than the open slather 5.8GHz favored by radio-clueless geek dreamers (contention, no range, no wall penetration).

Goldstein: In the US, the rules for 5.8 are quite generous, and potentially allow an unlicensed operator even more ERP than a licensed 3.65 GHz operator. There is no cap on ERP if it is "point to point", and as SkyPilot has demonstrated (28 watts ERP), you're legally "point to point" if you only transmit in one direction at any particular moment. SkyPilot is proprietary technology, but I suppose a 5.8 GHz high-ERP WiMAX system could be developed too.

There is a little bit of 3.65 GHz licensed WiMAX gear on the market, but it's costly. That band is available in some parts of the country (not most of the coastal regions, but in most of the interior) on a non-exclusive licensed basis. We have no 3.5 GHz band. The TV white space rules are extremely restrictive too, again very low power unless you have an outdoor antenna on a 12-meter pole, in which case you have merely low power. Sprint/Clearwire has cornered the 2.5 GHz licensed band, leasing most of the available educational licenses and buying up the commercial ones. The pattern's pretty clear; the FCC has helped the incumbent CMRS operators maintain their license value, intentionally crowding potential competitors out of the spectrum.

Harrowell: Of course, once you have the fibre (or high capacity point to point wireless) backhaul installed, then you can hang all kinds of access wires and radios on it, and I wouldn't be surprised if mobile operators were a major customer for muni-fibre; pretty much every network that's turned up **HSDPA** and beyond has seen their backhaul requirements erupt, and quite a few have started pulling fibre to their Node-Bs. [**Editor**: **High-Speed Downlink Packet Access** (HSDPA) is a 3G (third generation) mobile telephony communications protocol in the High-Speed Packet Access (HSPA) family.]

Goldstein: The non-ILEC mobile operators (T-Mobile and Sprint, plus the smaller ones) are desperate for non-ILEC backhaul. Sprint is putting in a lot of backhaul radios. They could be a user.

But the general idea of hybrid fiber/wireless needs more support, and I appreciate Mark's publicizing it. The USF model today is that rural high-cost support is for pure wireline networks, $50k+/home of Capex funded by REA and paid off by USF at $1k/month/home if that's what it costs to pull fiber to the farm. Wireless does not support 100 Mbps upstream, and has limited overall capacity for IPTV entertainment, but it sure beats dial-up, and it's a lot more cost-effective than rural glass.

Where I'm a bit uneasy about Mark's plan is that it reminds me of the FreePers' [Free Pressers] one too. Both are about creating a "third pipe" to compete with the closed duopoly, not about opening up the formerly-was-and-still-should-be common car-

rier plant to, uh, common carriage. Mark's emphasizes cost effectiveness, while FreeP's emphasizes big spending and high speed.

Does yer man at the Huff provide for Layer Zero openness?

The Policy Process and Canadian Examples

Cooper: **This is a policy process that is unfolding at a ridiculously rapid pace and I must speak a language that the policymakers understand. My goal is to crate an environment in which we can allow the best technologies to meet basic needs to prevail. You've identified lots of possibilities, none of which are contemplated by the incumbent communications companies in a meaningful way. We use the word WIFI/WIMAX as names that policymakers and the public may actually recognize, but they stand for a wireless solution to the basic service for the digital age.**

Once the local/middle mile backbone is built and out of the control of the incumbents, the superior last mile technologies will have a much better chance of prevailing. If we can use a community-based approach to basic service and get to 10X10 mbits for mobile and those who don't have a need for or cannot afford a the 100 mbits service, we will have liberated the communications network from the tyranny of the incumbents.

St. Arnaud: The business model that Mark Cooper advocates has been deployed in Canada in several jurisdictions with mixed results:

(a) Alberta SuperNet province wide network (b) Prince Edward Island province wide network (c) Villages Branches in Quebec

The model seems to work reasonably well in rural areas and small communities, where there is little or no competition from incumbents, but has failed to gain any traction in large urban areas.

Harold Feld: I think it's important to ask what we mean as the proper measure of success for these initiatives. For example [you wrote]:

St. Arnaud: The Alberta SuperNet is a good example. This was a province wide initiative to deploy fiber and WiMAX to all communities and public building and provide high speed Internet to all citizens. It was successful in reaching public institutions and using schools as hubs for WiMAX towers etc. But prices from commercial providers declined much faster than from the government sponsored program and as a result it quickly became obsolete for home use. Even schools are finding they can get much better pricing from local commercial providers (even in small communities) they can get from the government sponsored program.

Feld: Is the decline in cost perhaps related to the availability of the Alberta SuperNet as an alternative? I do not know the local data, but how do prices for availability compare in places where there is no government competition?

One of the impacts of muni broadband initiatives here in the U.S. was that it prompted private providers to deploy swifter, provide better service, and offer it at lower cost. If that proved the case in Alberta, you need to take that into account when considering whether the program failed or succeeded.

I have said many times that I do not expect public initiatives to out-Verizon Verizon. Verizon is much better at being Verizon than any non-profit driven entity will ever be. But Verizon provides a particular kind of service, oriented to a particular set of

customers, in a manner which maximizes its own profits. That's not a bad thing. But it provides us as citizens (or even as businesses) with a very limited set of options.

The community hotspot model offers a lot of valuable features that are not inherent in provision by private, profit-maximizing firms. It offers a safety net. It reenforces community institutions and potentially becomes a means of creating and delivering local content (we have some very good examples from community wireless networks). While it may do some things less well, like tech support, it will probably offer things that private carriers don't want to offer because the profit margin is much smaller -- like a plain vanilla connection without all manner of expensive fancy bells and whistles grafted on that drive up price.

I'm not arguing with your empirical observations or with your proper caution that we need to learn from existing projects when we consider a proper cost/benefit analysis. I am only suggesting we need to make sure we have appropriate and agreed upon metrics to determine success or failure.

St Arnaud: Harold: Both you and Mark, in a private exchange, have a brought up a very good point.

In the absence of projects like SuperNet, would customers, especially in rural areas, be getting the competitive pricing they are experiencing today?

The incumbents are so determined to prove that municipal networks will not work, that they will go to any length to undermine them. So it poses and interesting strategy - building out public funded networks as a deliberate strategy to push incumbents to more aggressively roll out their networks. Even if the public funded network fails, it still may be a better stimulus investment than handing the money directly to the incumbents - who will just pocket the money and do the minimum amount possible.

Earlier **St Arnaud**: The PEI network is a similar model, but has only recently been deployed

The most successful model was the Quebec Villages Branches where the government funded private sector companies to build condominium fiber networks in partnership with communities and schools. This allowed many facilities based competitors to serve a community e.g. a customer could have a choice of several Wi-Max providers. This ensured competition and drove down prices.

Joost van der Vlueten: Apparently there are two domains of BB-deployment: 1. areas where the market can do its work, given appropriate market regulation (e.g. unbundling, open access etc.): at least big cities and densely populated areas, where government (national or local) has a function mainly to prevent problems with parties that gain considerable market power, and can contribute to for example, FttH by aggregation of demand and other types of 'soft policy') 2. rural/remote areas, where a more active role of government (national or local, possibly also financially) is required.

Cooper: To be sure. Different technologies and regulatory models will be used in different areas**. The problem is that even in the large cities, where a "competitive" model is plausible, the cozy duopoly of cable and teclos are not giving us the network we need at a price we can afford.**

I recently did a debate in which I declared that 1968 was the greatest year in American regulatory history because in that year the FCC made the Carterphone and Computer Inquiry decisions.

These decisions liberated the communications network from the dead hand of the monopoly network owner. We have lost that framework and won't recover it any time soon.

Joost van der Vlueten: The art of broadband deployment is

A. to look at what is possible and to envision it as if it were there
B. do at little as possible and as much as necessary to get there.
C. take into account the interests and power to block of parties with vested interests in the situation-as-it-is.

The idealism required for seeing "A" functions as a blind fold for "C" Alas.

Mark Cooper: **The first draft of the broadband bill is out and it contains only $6.5 billion for rural areas. This is a massive victory, since we have stopped them from spending billions to further entrench the incumbents, especially if we can liberate some of the money for a community-based wireless approach, which is vastly superior in rural areas**.

http://www.cq.com/flatfiles/editorialFiles/budgetTracker/reference/docs/20090115stim-hsummary.pdf

Harold Feld then wrote at http://www.wetmachine.com/totsf/item/1442

After Blair Levin's warning to the world (and the financial markets in particular) that the stimulus package will not try to solve the broadband problems in this country and that people needed to stop dreaming in the tens or even hundreds of billions for broadband, no one should be surprised at today's announcement that the Administration/House proposal budgets $6 Billion for broadband primarily in the form of grants. Thank God!

There's an old Jewish joke about how a Frenchman, a Pole, and Jew saved Napoleon's life. Napoleon asks what they want as a reward. The Frenchman says his family were aristocrats before the revolution and he wants his family lands restored. "Granted," says the Emperor. The Pole says he wants Poland liberated and her pre-partition borders restored. "Granted," says Napoleon. The Jew says: "I want a real nice piece of herring."

Napoleon stares, turns in disgust to one of his attendants, and says "get this man a nice piece of herring from the kitchen and then get him out of my sight."

The Frenchman and the Pole turn to the Jew and laugh "You could have asked for anything! You idiot, that's the Emperor of France! And you asked for a nice piece of herring!"

"Ha," answered the Jew. "You think you're so smart? I'm actually gonna get my herring."

That's about how I feel about the broadband stimulus package. **Sure, I'd love to have had the feds build fiber out to every home. But I always knew that wouldn't happen. Worse, I figured that any HUGE pot of money would invariably end up chock full of goodies for incumbents with zippo oversight.** The old tried and true formula of putting money out like kids put out milk and cookies for Santa in the hopes that jolly old St. Seidenberg will festoon us with fiber, which has worked for us soooooo well in the past.

But a reasonable set of grant proposals, properly targeted, can do a boatload of good. Consider Mark Cooper's community hotspot approach, for example, or the work of ongoing projects such as the Mountain Area Information Network in rural North Carolina or the Lawndale Community Wireless Network in Chicago or any of thousands of projects in hundreds of communities working to bridge

the gap between connectivity and digital exclusion.

Certainly we have a lot of work to do to make this work for us rather than have it get sucked up by a bunch of parasites with great PR or artificially inhibited so that it can't "compete with the private sector" (i.e., can't actually provide any service or do squat except as a subsidy for incumbents). Just take a gander over at Universal Service Fund for what happens when you set up a program to give out money and then let the incumbents have a hand in structuring it. The folks at Free Press are already laying down a good marker of asking for public service and accountability conditions to go with the grants.
http://www.freepress.net/node/47297

I'll add that Waxman took this approach for HR 7000 last year. You want USF money for wireless? Than you and any parent or affiliate will provide reasonable roaming. Don't like it? Don't take money from the public trough.
http://www.wetmachine.com/totsf/item/1343 Same thing here. You want a taste of that broadband deployment money? Then you make sure we can tell how you are spending the public money so we can see if we get our money's worth, and you operate your network in a way that promotes our policy of an open and transparent internet. Don't want to do that? Then don't take the money.

I'll also add some recommendations for who should administer these grants. The Technology Opportunities Program (TOP) at Commerce has a long track record of funding programs intelligently that produced major bang for the buck. USDA's broadband grant and loan program also did a nice job - including creating competitive providers until the incumbents had them neutered.

But I'd also like to put in a plug for the Office of University Partnerships over at Housing and Urban Development, especially the Community Outreach Partnership Center (COPC) program, http://www.oup.org/programs/aboutCOPC.asp which gave grants for universities and colleges to partner with local communities (agency report/promotional material here). Granted I am biased because the program was developed by my Mom, Dr. Marcia Marker Feld (Hi Mom!). But funneling stuff through HUD and Ag as well as through NTIA and other usual suspects such as NSF will make this a "two-fer" or possibly a "three-or-more-fer" by pushing multiple agendas at once. For example, funneling money through HUD - especially if it is through COPC - not only helps broadband deployment to under-served communities, it furthers the urban agenda, gives a boost to higher education, and hopefully moves it away from places where the parasites hoping to suck up broadband money for no return have been making themselves cozy for far too long. [**Editor**: a useful example of trans-sectoral thinking.]

And as Blair Levin made clear yesterday, this isn't the end of the Administration's commitment to broadband. We have a lot of stimulatin' stuff we can do without spending more money - such as passing the Community Broadband Act to let local governments provide broadband service and freeing up more spectrum for unlicensed use.

So no bitching because the Administration didn't want to spend $200 billion on fiber to the home. There are no short cuts in public policy, and, if we want universal affordable broadband, we will need to fight for it the hard way, not get it as manna from Heaven. In the meantime, let's make sure those bastards don't steal our herring. Stay tuned . . . .

Paul Budde: We have similar examples here in Australia and New Zealand. As soon as an alternative operator launched a large scale wire-

less broadband network the incumbents dropped their ADSL prices and the bus models of the WiMAX operators collapsed. Despite 4 years of operation these operators (Unwired in Australia and Woosh in NZ) have not been able to even reach half a percentage of bb penetration. Also here in small regional pockets WiMAX is successful but if you count all of these bb subs together again you again wouldn't reach 0.5% of penetration. In comparison with the launch of mobile broadband a year ago (by the incumbent), this level of penetration amongst bb users is now over 20%.

Going back to Mark's points. I fully agree that once we do have a good backbone in place (based on open access) the last mile will find its own solutions and yes wireless could play a key role here, but once the incumbents finally realise that they will have to come to the party they - in most situations - will use their own last mile connections at prices that make it impossible for wireless operators to compete. The natural infrastructure monopolies will always make it nearly impossible to compete on infrastructure. But one could argue that if we have forced the incumbents to play the open network/ competition game and we do end up with low prices, open networks and competition between services (not infrastructure) that we have reached our goals. In the end infrastructure is a utility and I don't have a real problem to leave that in the hands of the incumbents (once again based on open networks, etc).

The Winds of Change Are Blowing

Editor: Finally on January 22, we had from Harold Feld an eloquent call to arms as the result of an exchange with Bruce Kushnick.

Feld: They have major transparency mechanisms, beef up the budget of GAO and the Inspector Generals, and created a new oversight Board. It will, however, require people to actually care enough to use these tools.

Kushnick: It will never have oversight. History, even in the good years, shows that it requires real work and NO ONE WILL TAKE ON THE INCUMBENTS! It ain't gonna happen...

Feld: Then why do you waste your time?

As for me. Well, as I've said before -- Ive had so many people tell me I'll never get what I'm working for that i don't usually listen when they explain to me why my getting it doesn't matter.

Kushnick: And the real problem is the state commissions, who are now essentially working for the Bells in many states, with some exceptions...

Feld: Perhaps you miss the logic of your statement above? You either get people motivated to care -- or you don't. Final tip -- people will care much more in hard times than in good times.

Kushnick: And one other caveat -- We're finishing up a report on CA phone bill prices -- all services, all carriers

Feld: But why if you don't think it matters?

Kushnick Enforcement is a joke. Transparency is a joke. And while I'd like to be optimistic like Harold, the only serious solution is way to radical -- divestiture/ structural separation -- so that the companies can't control the agenda of broadband deployment. ---

Feld: Hmmm....taste the bitter.

My brother the educator likes to say: "Parents are not keeping their "good" children at home. The ones they send to the classroom are the ones they have, so you better figure out how to teach them."

The people we have are the ones out there. If we can't figure out how to motivate 'em, we can give up now. If you've been working your ass off for ten years and nothing's changed, it m ight be because there's no way to win. Or it might be time to think about switching tactics.

Kushnick: I'd like to be a believer, but actions speak louder than anything with Susan's comments, it's clear no one is watching now....

Feld: And so it shall be, and so it shall always remain? If I am delusional, I must say it has proven a rather functional delusion.

Things I have heard in the last few years:

"You will never get the FCC to take on cable."

"You will never get any conditions on the 700 MHz auction."

"Getting rid of the media ownership limits is a done deal."

"Why on Earth are you bringing a complaint against Comcast? Even Kevin Martin won't want to do anything on network neutrality."

"You will never get white spaces, the broadcasters are too powerful."

"There is no way you can keep 3.65 GHz from going to licensed WiMax, Intel is pushing for it."

Of course, the same folks have always rushed in to explain why actually disproving each one of those statements didn't matter. Because the FCC taking on cable was "really" just Martin working for the Bells. And getting conditions on the 700 MHz auction didn't really matter because Verizon got the spectrum and they'll figure out how to cheat us. And old media doesn't matter and cross-ownership got relaxed so who cares that you saved all the rest. And the result in the Comcast complaint didn't matter, and it's bad anyway because it derails us from real rules and blah blah blah blah.

Oh we have lost ground too, no mistake. But after 700,000 people were willing to file comments in the 700 MHz auction proceeding on something as wonky as wireless Carterfone, I stopped listening to how the public can't get it.

There is a tide in the affairs of men, some shakey guy said. And also a good deal of inertia, I will add. Like the Hobbits of the Shire, we lived so well so long most of us forgot that getting justice and changing attitudes is a long fight measured in *years* -- against a well financed opposition with infinite patience and operating on multiple levels. Small wonder that, as years passed and people were prosperous, that the tide flowed with the incumbents and the few reformers and opponents found it hard going indeed.

But that tide is turning and the Shire is rousing. There is an interest and an energy directed at the management of policy not seen in far too long. It is for the most part still unformed -- more an allergic reaction to the last 30 years of free market triumphalism than an organized movement. But it is a real current looking for direction, and it lies with us whether we shall ride it and shape it or whether we shall allow others to but rocks and shoals in its path.

But our worst enemy is not the Bells or cable or any other company or incumbent. It is ourselves, unable to grasp that the world is changing and the moment for action and organization has come. Like the Children of Israel taken from Egypt, we may lack the essential character to be a free people even when the Almighty hands it to us on a silver plater. Go read your Bible and you will find that no sooner had God split the Red Sea and drowned the Egyptians that the Children of Is-

rael were bitching and moaning that it was too hot, they were too tired, how long was this gonna take, and why the heck did you bring us out of slavery to die in this miserable desert anyway? When the spies returned from the Holy Land, they reported "it is as a land flowing with milk and honey -- but the inhabitants are giants. AS GRASSHOPPERS WE FELT OURSELVES TO BE, AND SO WE MUST HAVE APPEARED TO THEM."

Well, if I am a grasshopper, I mean to bring a whole swarm of my friends with me to the party. But I do not think of myself as a grasshopper, or the opponents as unbeatable giants. It's a challenge, no doubt about it. And much unfairness and waste will still happen, bad guys will still manage to lay their cuckoo eggs in our public nest, pushing out productive projects and demanding we feed them. But we can also accomplish much good.

Unless, of course, you figure yourself to be a grasshopper. Look at yourself, Bruce. The stimulus bill is proposing tools that will make the job you are doing now ten times easier. But you are reflexively spitting on them. "Oh no I'm not, I'm just being realistic, this is good but it doesn't really change anything because nothing can really be changed because the way it was is the way it will always be."

So you take your reality, and I'll take mine. Feel free to pity me my naivete and to feel sorry for the inevitable crushing of my spirit. As for me, I'll find some younger grasshoppers with an appetite and a willingness to try something new. I think I see some green fields over the Jordan River.

IPv4 Numbers to Become Transferable and Consequently Property

A Introductory and Interpretive Essay for the Symposium Discussion on Pages 30-63 of This Issue

A Tipping Point for the Internet?

Catching the precise moment of a tectonic shift in a global system as large and important as the Internet may be viewed as an exercise in the improbable. However, I point out in this summary that I think we are precisely in the midst of such a shift. The largest portion of this issue is approximately 20,000 words of discussion of the ramifications of the exhaustion of the remaining pool of routable IPv4 address blocks - (pp. 30-63). This article serves as an introduction to that detailed discussion.

The RIPE policy announcement of December 16, 2008 sets conditions by which RIPE members who have IPv4 address block assignments can reassign some or all of them to other members. Although the action, at first glance may seem trivial, it is happening only because the pool of assignable IPv4 numbers, a limited resource on which internet growth depends, is running out. As the discussion in this issue points out, this change in policy has ramifications of which almost all of those who depend on use of the internet are unaware. This shift in the permissible use of the internet's most basic economic resource will have profound consequences. It is tantamount to the adjustment of stresses deep with in the San Andreas fault. The new policies will begin to send tremors through the global system in ways that bode ill for the open and competitive internet we have known so far.

The Context

IPv4 numbers are the fundamental building blocs of the global internet. While people can "participate" in the direct provision of internet content offering and web hosting based on transitory assignments of IPv4 numbers, the assignment of IP numbers on a permanent basis from a Regional internet Registry (RIR) is the only way a business can enable itself to route its customer's traffic via an ASN number. Such capabilities establish such a business as one of about 30,000 independent providers of Internet service in the world.

With the beginning of the Internet in the 1980s, a handful of universities and large corporations were able to participate in the deployment of TCP/IP as a transparent overlay of carrier networks. The protocol was hardware independent in a way that that other networking protocols were not. Because of this independence, the university community and its technology partners were able to construct an internetwork of networks on a large and rapidly growing scale. US government policy under the leadership of the NSF enabled the NSFnet backbone to interconnect to foreign networks in the late 80s and early 90s. At the same time, with a loosening of acceptable use policy, the NSF enabled small scale dial up commercial providers to connect to university endpoints. The university based Internet began to morph into the commercial Internet.

With the advent of the commercial Internet as marked by the decommissioning of the NSFNet backbone on April

1, 1995, the number of service providers, by then in the hundreds, quickly increased to thousands of independent ISP businesses. To facilitate the process of IP number assignments (the street addresses for the delivery of 'packets'), regional internet registries (RIR)s were established. In turn these registries were used by their members to administer the policies that they, the members, established for the gradual assignment of the IPv4 blocks of numbers needed by newly formed ISP business that wanted to become independent economic participants in the rapidly growing protocol overlay being constructed with a diverse mix of technology inputs including campus and corporate LANs, new greenfield point-to-point facilities, and fractional infrastructure products bought or leased form incumbent carriers.

The IPv4 numbers were simply an indispensable part of the TCP/IP protocol. To set up a service using TCP/IP, the service providers needed to be able to assign their customers unique IPv4 numbers from which and to which packets could be sent. The use of the IPv4 address blocks of specific sizes was supported on the basis that the ISP needed that many IP addresses -- no more, no less -- in order to connect useful things to the Internet, and that particular allocation was justified for as long as that need remained -- no more, no less. If the business was disbanded, the IPv4 blocs had to be returned to the registry to be reassigned. The blocks were not owned, They were not property. In economic terms, they were in effect "inalienable" -- not subject to being sold or transferred by the original allocation recipient to a third party, or to being purchased or acquired in other ways by a third party.

The blocks were like spectrum frequency assignments before frequency was auctioned. While spectrum frequency was there to be used according to the rules of the regulator, the IPv4 blocs were there to be used according to the rules set by the ISP members of the RIR. In the case of radio you broadcast on the frequency. Before auctions you could not claim to own the frequency. In the case of the ISP, IPv4 numbers are obtained directly from the reserve address resource pool administered on behalf of the ISP community by the Regional internet Registry, or alternately from a RIR member ISP. This hierarchical arrangement, with neutral RIRs at the top, competing ISPs at the next level, and individual users below that, parallels the organization of the banking sector, with a central bank (or occasionally, a "banker's club") at the top, competing lending institutions at the next level, and aspiring borrowers below, according to Tom Vest.

Vest suggests that this symmetry is no coincidence, but rather a product of the fact that IP addresses performs the same kind of "medium of exchange" function that money plays in the conventional economy -- and that the uses of IP addresses are subject to the same kind of systemic risks that can render money useless in certain circumstances in the wider world, for example, in times of extreme inflation or deflation.

Consequently IPv4 number assignments had economic utility by virtue of the fact that they were an indispensable part of the TCP/IP protocol that could be used as a transparent overlay" technology across carrier networks. The overlay was transparent to the networks that did not distinguish data network protocols from voice. Inter networks could be built atop various (telco) inputs, generally without requesting or securing their explicit permission about how the inputs would be used. Why? Because TCP/IP was just one of many data protocols to which the telcos were obligated to provide common carriage.

For the 1980s all this worked

well. Nevertheless, as long as IP addresses are essential but scarce, the prospect of a single entity having control of IP addresses created an inescapable conflict of interest for competing service providers. Recognizing this fact, beginning in 1993, ISP communities started establishing quasi-independent central bank-like institutions -- "regional internet registries" op RIRs-- to administer the distribution of these critical resources for the purpose of connecting useful things to the Internet.

The Tectonic Shift

Until December 2008, IP addresses distributed via the RIR system could not be bought or sold; i.e., they lacked the quality of "alienability," which economists regard as an essential feature of private property. But this lack was itself a feature – not a bug. IPv4 could be used only by an entity that agreed to create value by actively contributing to the Internet system. IPv4's lack of any other kind of use value, and the conscious collective decision to ***prevent*** it from acquiring exchange value, were the primary causes and rationale for creation of the RIRs and the "needs-based" allocation regime in the first place.

However, IP addresses can also have an "exchange value" if they have become scarce; that is, if those that actually need IPv4 now to attach useful things to the Internet can only obtain addresses from someone else that needs them less, and thus may be persuaded to part with them for some consideration. This is the strategy that was, in effect, chosen by the RIPE community when an IPv4 resource transfer policy was approved on December 16, 2008.

It was not the only conceivable strategy. For example, if the successor IPv6 addressing format had been transparently adopted by most or all community members, the scarcity and intrinsic exchange value of IPv4 addressing would have disappeared. But that did not happen.

Moreover, the fact that IPv6 has been rejected has even more far-reaching implications. Making IP number resources "alienable" strips them of one of the critical features that previously made them (1) irrelevant and transparent to countries and national jurisdictions, and (2) effectively manageable through voluntary, "self-governance" mechanisms.

If and when IP numbers become alienable, economic substance doctrine will eventually dictate that they ARE property for all practical (and legal, and regulatory, and taxation) purposes, regardless of whether some people might want to claim otherwise. The reason for this conclusion is that, given the economic impact of these markets, there will be litigation. When litigation occurs that the transaction (address block transfer) must have a meaningful economic purpose to be legitimate or sustainable in a court of law. An affirmative answer means that the IP Block transfer is one that involves property (something of economic value to the possessor. Going down this road invites government involvement. Because when IPv4 assignments become property, the only entity that can tell you what to do or what not to do with your private property is the property rights guarantor, i.e., the government.

Given the new opportunity to "own" IP addresses, how are incumbent services providers -- and especially incumbent facilities-based carriers -- likely to respond? The outcome of another recent and relevant privatization initiative -- i.e., U.S. spectrum auctions -- suggests a likely, if chilling, scenario. Given the benefits of securing scarce resources for their own customers, but even more importantly of blocking any possibility of competitive bypass, one may assume **that**

incumbent territorial facilities owners will do everything they can to acquire all available IPv4 addresses. Success would empower them to become the unilateral arbiters of all uses of the TCP/IP overlay -- **they will be able to demand whatever share they want of any IP-based service that they permit, and absolutely preclude any services that they dislike.**

Even if this scenario seems too pessimistic, it is hard to avoid the conclusion that, after all the usable IPv4 blocks have been distributed (circa 2011), the inheritors of RIR-era IPv4 will literally possess the keys to the kingdom. They will stand to achieve and enjoy permanent market power simply by doing nothing. That could mean that the "open Internet" could be finished forever -- or at least that the next moment of openness may come only after some new technology is invented that makes it possible to bypass TCP/IP in the same way that the latter made it possible to bypass the arbitrary restrictions imposed by telco facility owners.

The prospects currently look grim. On 24 January 2009 the ARIN Advisory Council (AC), acting under the provisions of the ARIN Policy Development Process, recommended that the ARIN Board of Trustees adopt: Draft Policy 2007-14: Resource Review Process

http://www.arin.net/policy/proposals/2007_14.html

This policy document will put in play a process even less transparent than that created by the members of RIPE. The only apparent way out might be for the registries and IANA to step back from the brink and establish procedures by which new entrants could be very gradually allocated the remaining numbers.

No, the Internet will not disappear, it will however become much more expensive to use. It will also likely fragment and lose its ability to stimulate growth and innovation**. Once the property right is recognized,its current beneficiaries will never willingly give it up. So no IPv6, or any other successor addressing format or technology that might reduce the rights and values of the new IPv4 ownership class -- will ever be accepted by them**.

The Internet then becomes the newly entrenched PSTN. It will become the domain of its new feudal lords who own not only not only the most indispensable component (the IPv4 numbers) needed to implement the the overlay protocol but this time the physical network as well.

The call then becomes where is the next bypass technology?

Acknowledgment: Thanks to Tom Vest who gave valuable assistance in getting all this in focus, and clarified many historical and technical details. Editorial comments or interpretations, especially regarding specific RIRs, are my own

Symposium Discussion November 18 2008 - January 18 2009

Who Should Provision Telecommunications Infrastructure?

Government or Private Entities?

Editor's Introduction

With the dawn of a new government of competence and intelligence there is reason to hope. The period between the American election and the inauguration produced much useful conversation on the Economics of IP Networks list. This discussion covers two months since the previous issue was just devoted to the lightwaves in the Netherlands material.

Thinking About Infrastructure

On November 18 Chris **Savage**: So some telecom seminar somewhere should have the following as its exam question:

Some infrastructure networks are provided by private entities. These include most electrical transmission facilities, gas pipelines, and long-haul railroads. Some infrastructure networks are provided by the public. These include roads, most (but not all) water and sewer systems, and most subway systems. Historically telecommunications infrastructure has been provided by private entities, generally akin to electrical transmission facilities and railroads.

It has been suggested that this model is fully adequate for the 21st Century as the Internet and wireless communications become the dominant forms of communications in the United States and the world.

Discuss this suggestion. Do you agree or disagree? Why?

Joe Kelly: I don't have time to write a detailed exam replay, but if I did, I might argue that the question is somewhat flawed.

Telecommunications networks are no longer 'akin' to electrical transmission facilities and railroads. Fifty years ago, electrical transmission facilities carried a form of controlled power into homes and businesses, while railroads carried trains, which in turn carried people and goods, between fixed points.

Today, electrical transmission facilities carry a form of controlled power into homes and businesses, while railroads carry trains, which in turn carry people and goods, between fixed points.

In short, both utilities perform the same service they did 50 years ago (albeit more safely). We can do more with electricity in the home and office today, but that's not because the nature of electrical supply has changed ... but because new devices and appliances were invented.

Fifty years ago, Telecommunications networks carried voice calls into homes and businesses. Today, they carry infinitely more. They enable data and video, they enable real time financial trading, they allow us to shop and eat out without having to carry cash, they keep us safe on the roads and in the air (and on the railroads).

The key difference between telecommunications and the water, electricity, gas and rail and roads networks is they are true 'utilities'. The level of innovation in the network is what distinguishes communications from utilities.

Wagter: Joe, I'll have to disagree with your approach. We are just at the beginning

of a massive transformation of our power grid fueled by the changes in how we use and produce and exchange energy, A lot of innovation is expected. Nevertheless I would never use that fact as a reason to argue for the privatization the physical transmission network because it is no longer a utility when you innovate a lot.

Over here in the Netherlands the physical network is being separated from production, sales, consumption and so on. The physical transmission network is a public utility, the rest is left to the market. The same view is actively supported by the EC. We do not privatize our road network because there is a massive level of innovation in our cars, what we can do with it and so on.

Fibers are conduits for light. The best fiber does....nothing (but guide the light). Only when you start pushing light through the fibers, modulate the light, and convert it into information it becomes a telecommunications network. You can innovate as much as you want in pushing and modulating light, it does not affect the fiber at all.

The fiber is a utility.

Tom Vest to Savage: This formulation is too general, and begs the question, "adequate to what end?" Private provision seems to be fully adequate in Korea and Japan at the moment, at least for the ends that make sense to me personally. It also appears, to me at least, to be adequate in the EU member states where I work and/or regularly visit. It seems to be profoundly inadequate in some other other countries, including some OECD member economies.

In all of the markets mentioned above, Internet applications and optical fiber-based technologies more work or less the same. One may agree that they're more like roads, or less like utilities as Joe suggests, but whatever you think they are, they're pretty much the same no mater where you happen to observe them.

Something else other than gross ownership structure is determining whether, when, and where the facilities that we associate with communications services are fully adequate, or not.

Cowen: [we have a] problem of being able to explain the transport of Communications across the world in any terms other than pipe metaphors.

We have thought that freight transport in those brightly coloured containers loaded whether on ships or railways or road or air is a better metaphor.

Most importantly for jobs and growth and prosperity the importance and industrial scale of what is done, the effects that it has on worldwide commerce and industry is better communicated if freight transport is used; and the value of the service and the way it is supplied is better understood by those with less technology knowledge.

Those that that is provide the 'Digital Freight Trains' of our generation can also communicate their importance to society with higher impact if this metaphor is used.

Barron: I don't think the freight metaphor is the right one unfortunately. It connotes "content" and puts us squarely back in the old paradigm of centralized production (ie Hollywood) and consumers. While that model isn't going away tomorrow, all of the transport metaphors ultimately fail to convey the game-changing nature of the net.

Hassinger: Kevin - I agree, I think for the transport metaphors to work one has to imagine that all shipments are carried in identically shaped boxes that each are self-aware, knowing who sent them, what they contain and where they are going. That way, any firm or individual

could grab one or more shipping boxes and deliver them. Some might set up huge UPS scale air-based operations, others might build larger standardized containers for moving massive numbers of smaller boxes on ships, others might just jump on their bike, drop by a depot on the way to work, grab a single box and drop it off on their way. Only that level of hypothetical model approximates the dynamics of the modern communications landscape, from backbone peering down to p2p voip and file sharing... Coincidentally, it would also make for a really cool global distribution network. =-)

Cole: I have my own difficulties with Eric Raymond's *"The Cathedral and the Bazaar"* that I have discussed with him (I think he mischaracterizes how cathedrals are actually built), but I think some of what he says about the Bazaar may apply here.

If we think of the "Internet" as a world-wide bazaar, with everyone allowed to be both buyer and seller (or giver and receiver), bringing to it whatever they choose to bring (within some legal limits), I think we capture more of the flavor of the constantly changing "content" and the constantly changing "links." Perhaps a "flea market" or swap meet would be even better, as those typically have participants both large and small, commercial and non-commercial, etc. and are often sponsored/hosted by some organization that provides the physical facilities -- usually paid for by charging exhibitors and attendees a flat fee, NOT a percentage of the action.

The world-wide swap meet is my current favorite, since it suggests both commercial and non-commercial interactions.

Barron: world-wide swap meet - i love it. Although with one crucial distinction: we are moving from the exchange of goods to the exchange of ideas, which is not simply a substitution. Physical goods do not have the propensity to alter each other (at least not directly and immediately) in the way that ideas or communications do.

So back to Chris' original question - the infrastructure (as distinct from the services) needs to be treated as a public good in order to maximize availability and create a level playing field. This is not a question of monopoly (government or otherwise), but merely an understanding that we are talking about public ownership as in beaches and the air we breathe. On the other hand, services should be open to all providers in order to allow innovation.

Cecil: One word, for me, covers it: functionality. That's what people want. What they do w/ it is pretty much limited to their creativity, but, more importantly, is always relative to whatever it is they are choosing to be at that moment in time. In that regard, I think Cathedral & Bazaar really miss it as do goods, ideas, or services. It may or many not be any or all or none of those. And therein lies the power of the Internet: functionality.

Barron: True - functionality is it. But how do you maximize it? We made do with 300 bps modems back in the day and are still making do with "dsl modems**". As long as we see the net as a series of tubes for different flavors of content, we will continue sucking for air. We can't simply say it's about whatever we dream up (however much that may be true). Perhaps if we talk about connectivity rather than transport, we get closer to heart of the matter.**

Cecil: **Functionality is neither minutes nor capacity. It is technology agnostic. Rather it is simply what people do with available resources. And, as you point out, more rational use and allocation of resources makes for better functionality. Lots of functionality rocks; fighting**

over minutes, pipes, subsidies, universal service for this silo or that silo, by contrast, sucks enormous amounts of air, among other things, but emits a few things too.

But functionality is broader than the Internet, and the Internet is broader than communications networks on steroids. On a planet where we consume 2 planets worth of resources (or more) we MUST do MORE with LESS. Functionality is about increasing the basic intelligence, adaptability, utility, and efficiency of planetary infrastructures upon civilization runs. Evolution says escape the limits of your system or die. It is very technology agnostic.

Whatever we do, therefore, let's begin with the end in mind, but let that end be broad, flexible, and adaptive: let it be a New Beginning(tm) rather than The End (tm).

The unifying goal, it seems to me, is simple: a sustainable living planet. Functionality helps. A lot. The rest, IMVHO, are details.

Fiber Unbundling in the Netherlands

KPN/Reggefiber Offer Fiber for Twelve Euros per Month

Nov 24 ***COOK Report***: OPTA cheers fiber investments" http://opta.nl/asp/publicaties/document.asp?id=2781

"OPTA publishes new fiber rules, KPN satisfied" http://webwereld.nl/articles/53704/opta-publiceert-nieuwe-glasvezelregels--kpn-tevreden.html

Budde: The crunch is that OPTA agreed to accept prices set by Reggefiber (41% owned by incumbent KPN). Fiber investment access costs vary between E775 and E 1025 per home. This results in monthly fiber access charges of between E12 and E15 and in some high cost areas they might go up to E14.50 and E17.50

Van der Berg: OPTA hasn't agreed yet ;-) it's a proposal. MONDAY, 24 NOVEMBER 2008 KPN/Reggefiber offer fiber for €12 euro/month

The Dutch will get wholesale fibre for between €12 and €17,50 a month. OPTA has just released its opinion on investments in fibre networks where a party with significant market power is involved. You have until January to voice your objections and then the digging starts. That's what tonight's published document says (more or less). To find out much more about this proposed deal and its background and how KPN and Reggefiber fit together see these ubervalued , super-duper, hi-James-we-copy-your-style fellow-bloggers and my own stuff. I hope those guys will go into comparisons on price etc as I'm now going to go to sleep.

To start with the sweet stuff in the back of the document. In the annex you can find the prices for the Optical Distribution Frame service of Reggefiber as the passive operator of the FTTH-network. The offer is for a wholesale service. The annex is written as the contract that operators of the active services on the network will have to sign. All costs are excluding VAT (19%) and applicable discounts

.

Depending on the area type there will be different CAPEX. The monthly fee per customer is dependent upon the area type and CAPEX. Normally the price should be at the first price, but it may go up by as much as €2,50 if certain penetration rates aren't reached to satisfaction etc. If an area has a higher or lower CAPEX, every 50 Euro's is another bracket and another 75 cents. There is a discount for bringing in more customers in a specific area. The discounts are as follows: **[snip]**.

This sounds like a good teaser for KPN as it will give them discounts in their retail side when they bring customers over from the DSL offers to the new environment. There are some connection, disconnection fees, Pop-fees etc. Most are one-off, some monthly. Energy is outside the scope. In the accompanying letter OPTA says it agrees with the Reggefiber's pricing. It will formally say so on December 19th and explain its reasoning after which it will notify Brussels and give you time to voice your objections. The Dutch Competition Authority NMA will start this week on a market test.

The main bit of the letter however is the market consultation of the proposed policy rules for price regulation for unbundled access to fiber optic networks. You have until December 8th to have an opinion. This bit is full of legalese. So I'll try to condense it. Everything is done in conjunction with the Proposed

decision for the unbundled access to the wholesale level of November 5th. In this proposed decision KPN is designated as a party with Significant Market Power.

OPTA nicely adheres to the stated goals of infrastructure competition. On paper it looks good, but realistically we won't see a second, third and fifth party rolling out a FTTH network under these conditions. The Netherlands is stuck with Cable vs KPN forever and wireless will play a minor role too.

OPTA will regulate the offer done by Reggefiber and KPN as it would with a copper network. So there are protections against price squeeze, discriminatory pricing and excessive prices. Nothing on delaying tactics, which is a bit of a pity as I would think that KPN's All-IP plans have already sown Fear Uncertainty and Doubt in the investors community and this plan just ups the ante.

The regulation will be by creating a ceiling for the prices Reggefiber can ask for its wholesale offer. The calculation of the price is done by using a variable All-Risk Weighted Average Cost of Capital. The WACC consists of the standard WACC for the copper network, with on top of it: An extra fiber WACC as fiber is a risky investment by itself. This one is thought to be variable and decreasing over time as the investors know better what they got themselves into. a regulatory risk WACC, which is stable as regulators are notoriously unstable and investors want to be compensated.

If the internal rate of revenue goes over the WACC than OPTA will either lower the ceiling or Reggefiber will need to quikckly invest more into less profitable regions. OPTA will evaluate the ceilings every three years based on the entire business case. Nice thing is that the Reggefiber is allowed to increase prices in accordance with the official Dutch inflation numbers. (Now how does that work in a time of deflation). Unfortunately there are no numbers on the height of the WACC etc yet. OPTA does promise to be vigilant and guard that Reggefiber won't make excessive profits by misrepresenting potential penetration rates, Capex per region etc etc.

Rood: The crunch is a bit different. Reggefiber acquisition has not yet been approved by our M&A authorities. Therefore Reggefiber FTTH management has till today very little interaction with prospective buyer KPN. As long as they kept operating inside the business plan KPN agreed upon, KPN is learning what it means if you have lost the gas pedal to a new entrant, they do not know where their feet are and the regulator is thinking they are with their foot on the brake ...

Soon we will learn who is holding the steering wheel here, who is looking in the rear mirror attempting to decipher in which way they are driving and who is kicking and screaming.

According to the IPTV news site
http://www.iptv-news.com/content/view/2565/64/
KPN receives "sufficient clarity" on FTTH regulation, but no mass roll-out yet November 25, 2008 - Dutch telco KPN has stated that it has received sufficient clarity "for now" from telecoms regulator OPTA and the Dutch competition watchdog NMa on the regulation of access to fibre-to-the-home networks, and will continue its investments in FTTH as announced earlier this year. The company has submitted a proposal in partnership with Dutch fibre-optic provider Reggefiber to the NMa, requesting approval for their proposed joint venture for the rollout of FTTH networks. KPN reports that although it will continue with its investments as planned, it does not yet have the intention to roll-out FTTH on a large scale in the Netherlands. KPN and Reggefiber report that they will assess all plans based on timing and location, considering whether new investments can be "justified commercially", and tak-

ing into account the regulatory framework. KPN will proceed with its current FTTH pilots in five Dutch cities, and a commercial evaluation will be conducted in the first half of next year.

Van der Berg: Formal statement by KPN --Den Haag, 25-11-08 KPN and Reggefiber have received sufficient clarity for now from OPTA and the Dutch competition authority (NMa) on the regulation of access to Fiber-to-the-Home (FttH) networks. In light of this, KPN and Reggefiber have submitted a proposal including access tariffs to the NMa, requesting for approval of their proposed joint venture for the rollout of FttH networks.

KPN will continue its investments in FttH as announced earlier in 2008. However, KPN does not yet have the intention to roll out FttH on a large scale in the Netherlands. KPN and Reggefiber will assess all plans based on timing and location, considering whether new investments can be justified commercially and taking into account the regulatory framework. KPN will proceed with its current FttH pilots in five Dutch cities.
http://www.kpn.com/corporate/nl/pers/persber.htm?contentid=6690 A commercial evaluation will be conducted in the first half of 2009.

Vincent Dekker: I always tend to read statements word for word. As a journalist I know many PR-words usually have more than only the obvious meaning. They are used so that a company can later say: 'Ahh, did you think we meant A? No, no, we really, really meant B...'

KPN now says it does 'not yet' have the intention to roll out FTTH on a large scale. Not yet doesn't give us any certainty at all on what their intentions will be tomorrow... And 'intention' is a beautiful word too. Who will blame them if tomorrow the CEO says: 'Ohh, we did not intend to yesterday, but today all of a sudden we got this chance in a lifetime...'

The next sentence: 'KPN and Reggefiber will assess all plans based on timing and location, considering whether new investments can be justified commercially and taking into account the regulatory framework.' describes pretty much what every sensible company will and should always do. And the third and fourth sentences only say that they will not stop their pilots (no surprise here) and will evaluate these after the roll out in these cities (very sound idea indeed).

So unless KPN is putting out a statement that they will definitely not proceed with or decide to any roll out anywhere else before June 2009 (which they won't put out) I don't think they are putting the brakes on anything. They just want to make it look like that.

In the meanwhile, their joint venture with Reggefiber (Glashart) is signing up new fiber customers by the thousands in more then 50 cities and villages all across the country.See my map of The Netherlands here:
http://www.trouw.nl/nieuws/economie/article1905291.ece/Glasvezelnet_werft__snel_klanten_.html. With green dots: ftth roll out complete orange dots: roll out in progress blue dots: actively attracting customers; no roll out yet. Would be a big shame to keep those customers waiting unnecessarily long, wouldn't it?

Trouw 2009, op dit artikel berust copyright.

Bit Torrent Changes to Use of UDP in P2P

Goldstein: Those who feel that it should be a crime to treat any packets differently from any others should pay attention to what BitTorrent Inc. is up to.

Bittorrent declares war on VoIP, gamers http://www.theregister.co.uk/2008/12/01/richard_bennett_utorrent_udp/

(Yes, it's that Richard Bennett, who opposed the ortho-neuts at the FCC Freak Shows.) <quote> Gamers, VoIP and video conference users beware. The leading BitTorrent software authors have declared war on you - and any users wanting to wring high performance out of their networks. A key design change in the P2P application promises to make the headaches faced by ISPs so far look like a party game. So what's happened, and why does it matter?

Upset about Bell Canada's system for allocating bandwidth fairly among internet users, the developers of the uTorrent P2P application have decided to make the UDP protocol the default transport protocol for file transfers. BitTorrent implementations have long used UDP to exchange tracker information - the addresses of the computers where files could be found - **but the new release uses it in preference to TCP for the actual transfer of files. The implications of this change are enormous.**

Felten: Will it kill it for a user that's simultaneously using bit torrent and VoIP and/or gaming or will it affect all users of the network indiscriminately ?

Sorry if that's a stupid question, my tech skills are non-existent!

Goldstein: Not a stupid question. The problem is not focused on the user himself, but upstream. A user who is doing both Torrent and VoIP might indeed be impacted, but they have the ability to stop their own uploading. In general, that user's access network (cable modem segment, for instance) also probably won't be hurt too badly, because each DSL or cable modem has a throttled rate. However, it's possible that the downstream (towards the cable user) will be congested, if the CMTS doesn't rate-throttle users, and leaves it to the cable modem. (I don't know which end routinely caps the downstream.)

The biggest problem occurs mid-network. I hesitate to use the word "congestion" because people see the word and assume that it refers to "scarcity", and thus can be solved by throwing capacity at the problem, which is not correct. I'll instead refer to the "flow dynamics". The flow dynamics of the Internet are based on end-to-end rate control triggered by packet loss. A common place for packet loss to occur is at the boundary between two ISP networks, especially a public NAP. (Another common place is at the egress, when the backbone is fat but the destination is thinner or busier.) When the NAP receives more packets for a given destination than that destination can absorb, it discards packets. TCP sees this and slows down ("slow start"), and the network is stable.

I don't like streaming in general because UDP doesn't do this -- a TCP application will slow down when packets are dropped but UDP won't, so UDP ends up taking a disproportionate share. We live with this in VoIP since nowadays the flows are relatively small. It's an issue for video because the flows are relatively large. I point out that Sun's 1980s NFS V1 used UDP too, which was a mistake, but it

was stable because it was intended to only be used on Orange Hose Ethernet, which had physical layer flow control and MAC layer fair capacity allocation. Across a routed network, it could be hell, and they ended up adding their own slow-start-like flow dynamics to it (a face-saving wimp-out vs. simply adopting TCP as they should have).

So what Torrent is now doing is bypassing the flow dynamics of the Internet and simply blasting away, like a DDoS attack, towards the destination. It won't slow down when TCP applications do. So loss rates at the NAPs in particular will rise. It's thus a social disease, not one confined to its users.

This is precisely the type of thing, like spam, that ISPs have to deal with as it comes up, and which can't be managed reactively by regulators following APA rulemaking.

Harold Feld: My thanks to Fred for forwarding Richard's piece.
http://www.theregister.co.uk/2008/12/01/richard_bennett_utorrent_udp/

I am amused that Fred seems to regard this is somehow the fault of "those who feel that it should be a crime to treat any packets differently from any others." Richard's article suggests that this is BitTorrent, Inc's response to the CRTC decision that *allowed* ISPs to throttle. Or, in other words, BitTorrent "declared war" because Bell Canada got what Comcast didn't get, the right to treat traffic differently based on application type.

As one of those who fought to make what Comcast was doing "a crime," I can't say I'm surprised by BitTorrent's conduct in response to the CRTC decision. To the contrary, this was one of the predicted *responses* to allowing network operators unfettered discretion in managing traffic.

Marti: Yes, but it's still irresponsible for Bittorrent to make this the default, hitting all network operators equally, rather than something it falls back to when it detects monkey business.

Maybe the open-source implementations (now less popular than the proprietary uTorrent, but still a presence) will take the sensible route, a tit-for-tat strategy. Stick to fair and square TCP, straight out of the Boy Scout Handbook, unless you detect "quality of service" actions on the part of your ISP or your peer's. Then switch to non-standard UDP trickery, periodically trying TCP to see if the problem ISP has seen the light.

Feld: One of the arguments made in the Net Neutrality debate (and I think Richard Clark raised this at the Boston FCC hearing) was that permitting Comcast to manage its network in the fashion it wanted destroyed underlying confidence in the basic protocols that have made the internet function as an open, interoperable environment. Network operators will manage traffic, application designers will respond. If no one (such as the FCC) steps in to ensure that certain basic ground rules get observed, rational actors will seek to maximize their own interest without regard to the welfare of others or the longer-term health of the overall network. The idea that network operators were somehow uniquely positioned to play referee was, I always felt, one of the fatal fallacies of the folks opposing the Comcast decision.

So lap it up all you Libertarian techno-determinists! The free market Nirvana you craved is at hand, where each network operator shall do what pleases him to manage traffic, and each application and content provider shall respond as pleases them. I cannot say I am surprised, either by the market response or by the anti-regulatory crowd to figure _some_ way to blame this on network neutrality.

Felten: I would point you out to this little post I wrote a few weeks back: http://www.fiberevolution.com/2008/09/copyright-warfare.html

This is warfare. It's not ruled by moral considerations, but any more in this conflict should be made with the full knowledge of what your opponent might do in response and the consequences of that...

I'm not blaming either party, but clearly both are irresponsible...

Feld: It is in some ways useful to conceptualize it as war, as I occasionally do myself. But it is important to realize that this is misleading.

The critical thing from a public policy perspective is that certain responses in an unregulated market are highly probable. If one subscribes to the theory that the result from the unregulated market (to the extent such a beastie actually exists) is inherently the best result because Coase promised that the market will always reach the most efficient result, then fine. But if you actually want certain outcomes, then one needs to accept that "the market" will not provide them without regulation.

It has been my perhaps biased observation that many (particularly in the engineering community) who oppose network neutrality regulation actually *want* a neutral internet -- with the exception of some applications or customers they view as "bad actors" who abuse the system to the detriment of the remaining users and network operators. The assumption seems to be that network operators want to provide a neutral internet -- it being assumed that it is in their best interest to do so -- and therefore will only act to prevent bad actors from undermining the system and will otherwise maintain the preferred internet model.

There are also some unspoken assumptions that network operators, as the heroes providing this miracle of access, deserve our respect and trust as the proper stewards of the internet (it helps that the engineers that are presumed to be making these decisions resemble those making these arguments). Another unspoken assumption is that since government is run by technologically ignorant cretins subject to the control of "special interest" (who are somehow not the carriers making the engineering decisions), government will inevitably bungle the job.

It's a nice picture. And it lends itself to little morality plays about wicked application developers, stupid greedy sinful bandwidth hogs, and dumb ass do-gooders like your humble servant who can't actually understand what we are talking about. While a fine basis for religious conviction, it is a lousy model for industrial policy.

In economic systems, it is rather foolish to talk about parties being irresponsible. It is also foolish to talk about bandwidth hogs or application developers or incumbent carriers as good or bad. Rather, what we should be doing is trying to predict their behavior -- preferably by looking at actual real world behavior to inform us about how people actually act in the market in question.

If you actually want a result, like a functional internet that conforms to today's existing expectations on openness, then you need to ask yourself how that will happen. You also need to ask what your error correction will be when it turns out you made the wrong choice, because I stone cold guarantee that any system involving human beings using imperfect information to predict a dynamic future will include wrong choices. And for all you Libertarians out there, failure to take action for fear of making the wrong

choice in an imperfect world is as much a choice in this context as the decision to impose a rule or enforce one. This whole "do no harm stuff" is simply code for "take no responsibility and pray it all works out" (with Plan B usually being "no matter what actually happened or what we previously predicted earlier, convince everyone it really _did_ all work out").

Effective public policy must be sufficiently robust to survive irresponsible actors. Irresponsible actors are an utterly predictable set of actors in a field of stakeholders this large. Blaming the actors for being irresponsible is as much a waste of time as blaming incumbents with market power for exercising market power absent regulation preventing it.

RIPE Approves Private Market for IP v4 Numbers

A Discussion of the Operational and Economic implications of the Depletion of the Internet's Currency

Economics of IPv4 Numbers

Editor's Introduction

[Please do not read this without having first read the general introduction on pages 15-18 above.] What follows is a very long discussion about the exhaustion of routable blocks of IPv4 addresses. It is a subject that the *COOK Report* has reported on more than once during the past year and a half. Its outcome seems ever more likely to have a profound impact on the economic functioning of the Internet. Players who will be impacted are for the most part unaware. Yet inexorably the problem is sneaking up and the formulation of policy designed to deal with it in such a way that the value of the internet as an open system with a low cost of entry has a greater chance of being preserved is being over looked.

Thinking about process does not come easily to the technologists who are making most of the decisions. Thoughts of markets do and despite the fact that 20 years of embrace of free markets have just ended in disaster in the US, a strong assumption in the discussion that follows is that once again we have nothing to be concerned about. I disagree but like Harold Feld can't claim to have **the** answer.

A Plea for Help in Understanding the Problem

On December 12, **Feld**: Against my better judgment, I am starting to look at the issue of IPv4 address exhaustion and other economic issues surrounding the IP address system and routing separate from all the blather about the name space.

Has anyone done any solid economic analysis on this? My gut tells me that the current system proliferated initially for a variety or reasons and then achieved critical mass necessary to have massive network effects.

Vest It had serious competitors right up to that point, especially OSI. At the last ARIN meeting, during a panel of graybeards on "What Would John (Postel) Do?" Paul Mockapetris et al. revealed that Jon and several others were fairly continuously/actively involved in the equivalent of "lobbying" for IPv4 (i.e., engaging in non-technical "competition by other means") until OSI ceased to be a credible challenger.
http://www.arin.net/meetings/minutes/ARIN_XXII/joint_transcript.html

Feld After about 1994 or so, no one was going to try to adopt an alternate system.

Vest: Beginning almost four years before that, the IETF community had recognized that IPv4 addressing and flat/non-aggregated routing would soon impose absolute limits on the number and size of entities directly able to participate in Internet routing. Without "enough" IP addresses, some users and services would be excluded, and as a result all of the existing IP-connected users and

other resources would be deprived of their (non) contributions. Without IP, such excluded agents would only be able to exchange traffic/value/whatever with other Internet-based economic agents through various bilateral, ad hoc, barter-like transactions; or eventually they might all band together and develop some other, non-IP based exchange mechanisms, in effect creating permanently partitioned exchange domains in which neither could become as useful as they might otherwise be together.

Beginning in the early 1990s, these kind of deflationary pressures outside of the US had already fueled demand for additional regional IP address distribution & registration centers (RIPE was established in 1993). Eventually, five were established, driven by the same dynamic -- or at least the same winning public rationale (a concession to my more cynical friends).

In addition to this "deflationary" risk, existing flat routing conventions created negative economies of scale in the costs and stability of inter-institutional routing. Initially, each new routing participant required its own (one) entry in the routing table, and before long rapidly changing telecom regulations, falling costs of telecom inputs, and proliferation of administratively independent routing domains created widespread opportunities for routing system participant to multi-home (and later, to engage in even more fine-grained varieties of traffic engineering). Each one that did so added to the multiplier for routing costs that has to be borne by all participants. The resulting negative externality operates very much like general price inflation in terms of who is affected first, whose interests are advanced vs. undermined, and what happens if it is allowed to go unchecked for too long.

In order to mitigate and manage (not eliminate) both systemic risks, CIDR and the RIR system of "needs based" address delegation was developed.

That solution didn't sit will with all stakeholders however. There have always been advocates for a purely decentralized, market based approach to address distribution. Their vision is spelled out in RFC 1744 (http://tools.ietf.org/html/rfc1744). The IETF PIARA group investigated these issues fairly extensively, and one artifact of that debate is Yakov Rekhter's paper/slides called "Address ownership considered fatal" (ftp://ftp.ietf.org/ietf-online-proceedings/95apr/area.and.wg.reports/ops/cidrd/cidrd.rekhter.slides.ps).

With the exhaustion of the unallocated IPv4 pool looming, and uncertainty about IPv6 actually increasing as a result, we are now living through the same arguments once again.

Feld: The failure of alternate roots in the mid-1990s to the present can be explained for many reasons, but I believe the primary reason is that the current system works reasonably well and there is very little value to be gained by any player in adopting an alternate root.

Vest: You may well be right, for both DNS and IP addressing, but as the above should make clear, there were a lot of potentially existential-scale risks facing TCP/IP in the early 1990s. If IP addressing was just another general input for Internet service delivery (e.g., like T1s vs. SONET, or Unix vs. specialized operating systems/platforms), it seems likely that we would have seen a gradual migration/replacement there too.

IP Numbers Different from Domain Names

Feld: But IP addresses are fundamentally different from names, as I am sure I do not have to explain to anyone on this list. The distribution of IP

addresses has real economic consequences.

Vest: The distribution *can* have real economic consequences, in the sense that given possession of all of the other inputs required to produce some Internet content or service, a lack of sufficient interoperable IP addressing can be an absolute bottleneck**. An absolute lack of either - infrastructure or IP address resources - can absolutely preclude productive use of the other for the purpose of creating Internet services. But no amount of IP addressing can substitute for (an absolute lack of) the other inputs necessary, and vice versa.**

The only "use value" an IP address conveys is the potential to easily attach the devices you use to create online content and services to the pre-existing collection of IP-mediated Internet resources, for the purpose of sending and receiving mutually comprehensible packetized traffic. To produce the greatest possible value or usefulness, infrastructure and IP number resources have to be maintained in some kind of productive ratio - not unlike the ratio between money and "real economic factors" that monetary policy authorities try to maintain (a.k.a. the "quantity theory of money").

Feld: The availability of IP addresses impacts the ability to communicate.

Vest: That's true, but only in the sense above. Availability of hardware, software, telecom inputs, and the existence/reachability of others possessing an interoperable platform and who are willing to interconnect are also essential in the same way.

Feld: But in addition, the actual distribution process has consequences for the global network as a whole. But I'm having a devil of a time trying to understand how it all fits together.

Vest: Hopefully this will help a little ;-)

Feld: Has anyone written anything exploring how these factors interrelate? I'm seeing bits and pieces, but it is awfully hard to pull together a whole picture. I'm not even sure I can get all the relevant factors assembled.

Vest: Well, Eliot Lear, William Lehr, and I put something together for this year's TPRC conference:

http://eyeconomics.com/backstage/References_files/Lehr-Vest-Lear-TPRC2008-080915.pdf

Jean Camp also wrote for the same panel:

http://tprcweb.com/files/IPv6%20in%20our%20Lifetime9_4.pdf

Geoff Huston recently wrote something for the Cisco IP journalJ (my rebuttals go out in the next two issues):

http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_11-3/113_ipv4.html

And Milton Mueller has done a few things reflecting his own distinctive view(s): (In 2005 he was advocating that the ITU step in compete against the RIRs; today he has concluded that government participation in the address distribution is abhorrent. The only thing that seems to remain consistent is his view that absolute personal privacy, for artificial as well as natural persons, trumps all other values, including the existence of a working Internet).

http://internetgovernance.org/pdf/igp-v6.pdf

http://tprcweb.com/files/RIRs-TPRC08-Mueller-Kuerbis.pdf

http://internetgovernance.org/pdf/RIRs-IGP-hyderabad.pdf

I have hundreds more references that I've been meaning to get up on my website Reference page. I'll let everyone know when that happens. In the mean time, please let me

know if you have questions, or if there are other ways that I can help...

Geoff Huston: [Written anything?] no, not at all, or at least not in any useful, clearly presented and well researched way in my humble opinion, other than a really impressive paper from the OECD in the middle of this year, which as far as I am concerned is about the first and quite possibly the last word in this area. see http://www.oecd.org/dataoecd/7/1/40605942.pdf

COOK Report: Does it look like IP Address transfer markets are going to carry the day?

Vest: Yes

COOK Report: I hope not..... but that is gut emotional feeling on my part not anything related to well rounded knowledge.

Vest: Knowledge or not, your instincts are correct.

COOK Report: Except i would think that after what just happened to the global economy we'd be a little slower in taking up the market solution.

Vest: As it turns out, Santayana was an optimist. Sometimes even those who do learn from history get to repeat it anyway :-\

Goldstein: I agree with Tom that it's happening, but also think it's a good thing.
They stopped handing out homesteads over a century ago, but there are still lots of farmers in the western US. Today's IPv4 addresses are in the homestead era (minus the genocide of the previous occupants, of course).

Vest: By the time they stopped handing out homesteads, land has ceased being a primary prerequisite for political participation, and had also stopped being the primary input for economic activity and capital formation. By the time they stopped handing out land, it had become fully substitutable for these purposes; there were a variety of other options available other than landholder vs. disenfranchised tenant (a.k.a., serf). The Internet could eventually (soon) come to encompass all media and communications. Unless you happen to believe media and communications functions are not very important, and political participation ins not all that important, and innovation is just as effective when one person can pursue their vision of the future as when 1,000 can pursue their visions, then I think you might want to reconsider your position.

Goldstein: What just happened to the global economy doesn't prove that markets are bad per se. It proves that deregulation is not good per se. Scoundrels and thieves (the kind who wear expensive suits) were engaging in wholesale brigandry on a scale never before imagined, and justified it by claiming that they were simply engaging in "markets". But a market only works when the boundaries of behavior are understood and when there is an authority to enforce contracts.

Vest: In a recent interview with *Vanity Fair*, Joseph Stiglitz had this to say on the subject:

"I had opposed repeal of Glass-Steagall. The proponents said, in effect, Trust us: we will create Chinese walls to make sure that the problems of the past do not recur. As an economist, I certainly possessed a healthy degree of trust, trust in the power of economic incentives to bend human behavior toward self-interest-toward short-term self-interest, at any rate, rather than Tocqueville's 'self interest rightly understood.'"

http://www.vanityfair.com/magazine/2009/01/stiglitz200901

If you sincerely want a working market, it might be a good idea to take into consideration the existing incentive structure, and try to devise mechanisms to encour-

age "efficient" or at least sustainable outcomes. If anything, the recent crisis has proven that "market" can no longer be treated like some kind of magic word that automatically legitimates whatever outcome one decides to apply it to.

Goldstein: What we're looking at in the IP world is a still-functioning market. As was demonstrated aptly by the failure of the Soviet economy, command economies don't work either.

Vest: An excellent comparison. Would you say that the "big bang" that Jeffrey Sachs and others vigorously championed for the post-Soviet transition was the best solution? Would you say that shock therapy resulted in the kind of "nice market" that you would like to see govern the administration of IP number resources?

Goldstein: Market forces are a natural lubricant that allocate resources pretty well, especially on a micro-economic level. They should be harnessed, not worshipped.

Vest: You were better with your first formulation; **market forces can be a natural lubricant that allocate resources pretty well *when boundaries of behavior are understood and when there is an authority to enforce contracts.* But even that is not enough to *guarantee* "pretty good" outcomes, especially when you start off with a heavily skewed distribution -- and especially when the commodity itself is an absolutely non-substitutable bottleneck input for a wide range of important industries.**

Goldstein: It's not as if we have a choice. This simply will happen. There is no crisis, no end of the world. Right now there's no address-block market because the price of new ones (homesteads) is so low, but once those are gone, the market will find a price. This will also encourage NAT, which will reduce demand, and help lead to what I've discussed before, more NAT-friendly applications. And that is a step towards a desirable decentralized model.

Vest: Maybe. It's also a step toward crisis and nationalization, or absolute technical breakdown. A clear-eyed view of how markets work in the real world would suggest that your preferred outcome is just one of many possible outcomes, and probably not the mostly likely one either.

Goldstein: I am more optimistic than you. "Don't worry, be happy." If nobody screws it up, things will work, from the bottom. Remember, IP addresses are not government property. They are simply network names, labels agreed to by the users. ISPs run the show. yes, there is a risk of fragmentation, but even that would not be the disaster that people fear. Relays could be set up.

Remember, I'm also positing that the current model, with a single flat centrally-controlled "address" space visible to the application, is wrong. But we discussed that last month. If address markets fail to do the job, then the DIF model could pop up from the bottom.

Editor: earlier **Vest** had written: The Internet could eventually (soon) come to encompass all media and communications. Unless you happen to believe media and communications functions are not very important, and political participation ins not all that important, and innovation is just as effective when one person can pursue their vision of the future as when 1,000 can pursue their visions, then I think you might want to reconsider your position.

Goldstein: Why? Land was resold. Land was inherited. Land was subdivided. Landholdings were consolidated. I'm not talking about land as a qualification for landed status; I'm talking about the stuff you grow food on.

Vest: In that case you're not talking about the kind of land that is relevant to IP addresses. To paraphrase something that Yakov Rekhter wrote in the previously cited slides, we may choose to ignore the fundamental properties of the system under discussion, but ignoring them doesn't change the facts. Once land ceased to possess those other economics-transcending features, it became just another investment option, with some people choosing to buy as a input for food production, and some people choosing to buy for speculative ("development") or anti-competitive purposes - and many choosing to make their livelihoods in some way that required no land at all. **Once a market is established, intentional distinctions like this become moot.**

Goldstein: I think you and Gordon are confusing me for a right-winger who uses "market" as a religious mantra. I'm an unabashed leftist, a believer in democratic socialism (the western European mixed-economy style) a la Mitterand. Regulation and state investment are fine with me. But you can't ignore markets, especially on a small scale, the little transactions that take place between individuals and productive businesses. That's not the same thing as the transactions that took place among I-banks.

Vest: I understand completely, you want a "nice market" -- i.e., the kind that enables transactions that seem reasonable to you, but which is not vulnerable to the kind of manipulations and excesses that have recently laid the financial sector low. But the *only* thing that ever delivers outcomes like that is market transparency -- i.e., the ability of outside parties to understand how the market works, who are the principal players on all sides, and what "prevailing prices" are. In general, these are precisely the kind of conditions that IP number privatization advocates have rejected, either on philosophical or practical grounds. Absence of transparency is the hallmark of Internet-related input markets, so it's absurd to imagine that some other outcome is likely to emerge "naturally."

Goldstein: If some guy named Vito from Belleville bought an insurance policy on Gordon's house worth five times its value, he'd have gone to jail. But if a guy at Lehman Brothers did it, he got a bonus. When he sold a policy on ten times its value to Goldman Sachs, the Masters of the Universe at both companies got a bonus. I think what they did was not really covered by "regulation" in the usual sense; it sounds more like crime. But you know the old cliche; politicians get more play out of fighting crime in the streets than crime in the suites.

Vest: It may sound like crime, and perhaps it may actually be criminalized hereafter -- but such a development will only come after the fact, after transparency has been created as s byproduct of the industry collapse/post-mortem/bailout.

Goldstein: [refernecing Tom Vest's quote of Stiglitz's *Vanity Fair* interview above] I agree with Stiglitz. There must be rules. Don't listen to scoundrels who say, "trust me".

Beware of Unintended Consequences of New Markets

Vest: The thrust of the IP number resource transfer proposals considered to date is "trust me."

Goldstein: However, the IETF fundies seem to think that addresses must never be bought and sold, must always be handed out by Gosplan, and by the way they're the new Gosplan, so bend over. Uh, no. They screwed up big time. They put on the faith-

based hat and preached IPv6, and it failed, so now it's somebody else's turn. Not another dictator. The answer is obvious and it is taking shape.

Vest: Okay, so it's obvious that you bear a lot of animosity to the established technical coordination institutions -- fine, perhaps they deserve every bit of it. But this just brings us back to the attitude that I used to describe as "Leninism," i.e., the view that we live in the worst or all possible worlds, and that any change -- even one precipitated by the total collapse of the existing system -- represents a politically useful step in the right direction. You should be a little careful about what you wish for: if you truly think that nothing could be worse than the status quo, you may just get a chance to find out.

All I'm suggesting -- all I've ever suggested -- is that transfer supporters jettison their rose-colored glasses and seriously address the most serious issues created by resource transfer markets -- very few of which are actually addressed by the OECD report.

Goldstein: I'll say it again: The Internet is a business model, not a protocol. Too many of the IAB/IETF/ICANN insiders have forgotten the model and have become a parody of the bad old ITU.

Vest: **Do business models need rules too? Where do they come from? How do you determine whether they're sustainable over time or not -- is it always necessary to find out the hard way, as we are today?**

Goldstein: What we're looking at in the IP world is a still-functioning market. As was demonstrated aptly by the failure of the Soviet economy, command economies don't work either.

Vest: An excellent comparison. Would you say that the "big bang" that Jeffrey Sachs and others vigorously championed for the post-Soviet transition was the best solution? Would you say that shock therapy resulted in the kind of "nice market" that you would like to see govern the administration of IP number resources?

Goldstein: Sachs was a jerk (I have a stronger word in mind, actually) too. His "shock therapy" hurt the populace, gave the wealth to a small handful of thieves and scoundrels, and unintentionally opened the way for the re-sovietization of the economy, as is being pursued by Putin. That by the way was not a true market either. The conditions for a market economy didn't exist. They just changed some names around.

Vest: I agree with your sentiment, but I question your reasoning. Every reason you cited to abhor shock therapy was an *unanticipated outcome* of what the vast majority of experts at the time thought was a reasonable plan. Remember, shock therapy was deemed a relative success in Chile, and also in Poland, before problems were encountered in Russia. **What was the big difference in Russia? Arguably, weakness of institutions and rule of law, and complete absence of transparency.**

Under current circumstances, an IP number resource transfer market would operate within very weak institutions, no well-defined law, and the complete absence of transparency.

Goldstein: There is no such thing needed for IP number resources. When somebody wanted a farm in Kansas, they could buy acreage from a homesteader. The Soviet economy was, in contrast, very much on a macro level, huge state enterprises turned over to well-connected oligarchs. A market has to start at the bottom. Small private business should have been encouraged first, with state ownership retained for the

giants, and shares sold off gradually. (China is more along that model. It seems to work, though it is far from ideal for all stakeholders.)

Vest: Maybe you are right, but that's not how it happened, in either Russia or China. And if resource transfers are approved in anything like the form suggested by current proposals, there won't be any "shares" to worry about; the factories will be transferred en bloc to the current factory managers.

That said, I think that this is a productive line of discussion. If we could just get some real facts to apply to these comparisons, then maybe it would be possible to understand the kind of circumstances under which transfers would not lead to near-tern industry collapse.

Goldstein: So there will be no IP shock therapy?

Vest: Oh I don't know, what's the actual definition of "shock therapy" -- or if you prefer, "big bang" style economic reforms? Simultaneous total transformation of several core economic systems, on the theory that the alternative ("sequencing") is less likely to succeed. What kind of core functions?

1. Distribution architecture - e.g., from centralized to decentralized.
2. Eligibility criteria - e.g., from categorical ("all citizens" or "need-based") to competitive (i.e., markets).
3. Medium of exchange -- e.g., from one currency to another, or in this case from "credible promises of new Internet production" to whatever considerations are acceptable to incumbent IPv4 dealers.

What are or were the theoretical justifications for shock therapy?

1. The current system has reached end-of-life, and nothing less than total transformation will make it better.
2. The incentives are misaligned, let's introduce the profit motive to get things moving in another direction.
3. The medium of exchange has failed -- usually that means hyperinflation, not severe deflation as in this case -- so let's dump it.
4. The locals are too insular, too protected -- let's introduce some foreign capital/ competition/expertise.

What are/were some of the more common effects of shock therapy?

1. Deflation -- but the therapeutic kind rather than the Great Depression kind, because the previous excess was hyperinflation. Unfortunately, we're already experiencing deflationary pressures with IP number resources because of the non-substitutability of IPv4.
2. Initial, severe disruption -- lots of industry restructuring, with lots of pain for those working in the affected industries.
3. "Insider privatization" -- a necessary and accepted evil in sectors where expertise was scarce, but which can lead to excesses (with names like "crony capitalism", "kleptocracy", etc.), esp. in situations where public transparency or property rights/rule of law is weak.
4. Speculation / Consolidation -- i.e., the opposite of competition, but good for getting prices up, which spurs investment!

And sometimes, much later:

5. Recovery, normalization, usually with PTSD -- those who lived through it wondering if there might have been an easier way...

Tell me, what sounds out of place here?

Goldstein: Some little ISP in Sheboygan, WI with a /16 may sell it to another ISP in Cheboygan, MI, in exchange for the use of a /20 subnet for his remaining customer base. A company with a /16 in Springfield, MA will sell it to a company in Springfield, MO and let its ISP assign it a

/24 block or two that is adequate for its needs. That's the sort of transaction that will happen.

Vest: These transactions sound so nice and idyllic. I sure hope that those little ISPs in Sheboygan will have more money to invest in IPv4 than the 2-3 largest US facilities-based incumbents, which will be scouring the Earth for any/all liquid IPv4 resources. I also hope that the Sheboyganites will also have deeper pockets, broader contacts, and more sophistication that the prominent network service companies that are now preparing to launch pre-emptive speculative resource grabs in order to capitalize on the certainty of ever-appreciating IPv4 prices that will accompany and accelerate with the final rejection of IPv6, at least for a decade or longer.

Goldstein: More private networks will move into the 10.x private space behind NATs. A brokerage business will open up. It might be wide open and transparent, if the IAB glitterati do not stand in the way.

Vest: So, why exactly would markets naturally favor "open and transparent" transfer arrangements now, when there is no historical precedent for that assumption/outcome in any other industry, and every precedent of provisioning/pricing in every other ISP service input market (colo, transport capacity, interconnection, hardware, etc., etc.) suggests the exact opposite? Surely you don't think that it's the IETF's fault that interconnection criteria are so opaque, or that bulk IP transport prices are so wildly variable (even for the same capacity increment over the same route at the same moment in time), depending on who you are?

Waclawsky: The IPv6 future revolves around two key points: 1) Is there anything you can do with IPv6 that you can't do with IPv4 and related technology? - does a good rational technology reason exist for its deployment? ....maybe that why IPv6 has been in "who cares" limbo.... And how about the opposite question? 2) Is extending the current network faster, cheaper, easier etc than any kind of IPv6 build out? ....a interesting corollary might asking: will IPv4 go away once IPv6 is "on the scene"? My 2 cents.

David Meyer: All true, but the argument has moved on. The question is no longer IPv4 v. IPv6. Its how are we going to get from the heavily NAT-PT'ed world of Carrier Grade NAT (CG) to something that resembles today's e2e Internet. One might also note that that the deployment of CGN and IPv6 are not independent. See, e.g.,

http://www.1-4-5.net/~dmm/talks/NANOG45/iteotwawki-nanog45.pdf

Huston: Nice pack, but beware of false dichotomies - its a big and creative world out there, so the end point may be well distanced from our original intention once we embark on making drastic changes to the network's architecture to extract further scaling leverage from IPv4. For a depressing commentary you might want to look over http://www.potaroo.net/presentations/2008-11-17-ipv6-failure.pdf (and thanks for your slide here Dave!)

For an more optimistic perspective there's http://www.potaroo.net/presentations/2008-12-10-ATNAC.pdf

Meyer: Nice deck. More thinking needed...

RIPE Adopts Allocation Policy December 16, 2008

Then on December 16 **Tom Vest**: Rudolf's previous announcement has just been made true.

[**Editor**: Rudolf van der Berg had posted a note back in

October suggesting that Ripe had then adopted the transfer policy outlined below, which was formally approved today.}

Vest: The Internet -- doors closing soon! Enjoy it now, while supplies last!

Filiz Yilmaz: PDP Number: 2007-08 Enabling Methods for Reallocation of IPv4 Resources

Dear Colleagues,

Consensus has been reached, and the proposal described in 2007-08, "Enabling Methods for Reallocation of IPv4 Resources" has been accepted by the RIPE community.

The related RIPE policy document is now updated, published and can be found at:

http://www.ripe.net/ripe/docs/ripe-441.html or http://www.ripe.net/ripe/docs/ipv4-policies.html

Further implementation details of this policy will follow soon. The proposal is now archived and can be found at:

http://www.ripe.net/ripe/policies/proposals/2007-08.html

COOK Report - added on January 25 2009 during process of editing for publication: *In one sense what RIPE has done is pushed the allocation decision out to a local Internet registry in other words the ISP. The ISP uses its assigned numbers to connect its customers. If an ISP has a bloc of numbers it is not using, it can now transfer them to another ISP within the RIPE region (presumably for money) and that ISP can use the bloc to connect new customers. The policy pushes the allocation process one level further done and leaves the details of the execution to the decisions of the members giving reallocation power to potentially 11,000 existing local registries that is to say - local ISPs. To participate in the reallocation market one must already have an assignment and an ASN number.*

The RIPE policy goes on to add "LIRs that receive a reallocation from another LIR cannot re-allocate complete or partial blocks of the same address space to another LIR within 24 months of receiving the re-allocation.

The RIPE NCC will record the change of allocation after the transfer. Please note that the LIR always remains responsible for the entire allocation it receives from the RIPE NCC until the transfer of address space to another LIR is completed or the address space is returned. The LIR must ensure that all policies are applied.

Re-allocated blocks will be signed to establish the current allocation owner.

Re-allocated blocks are no different from the allocations made directly by the RIPE NCC and so they must be used by the receiving LIR according to the policies described in this document."

The policy creates an entirely new level of internet authorities that must administer their routing allocations with the same care as the much larger authority.

St Arnaud: Is not the opposite [from Tom's assertion that the internet is closing] true? The RIPE re-allocation policy will **[Editor - could]** allow eventually 193,000 more allocations (about 17 times more than there currently are).

This policy is opening the doors for continued growth (but may be the death knell for IPv6)

Vest: I guess the future will shortly provide an empirical answer to that question.

David Conrad: Yep.

Vest: My own observation is that the existence of supplies does not always entail the appearance of "supply" in an economic sense.

Conrad: Indeed. For example, there is a large supply of plutonium 238 but for the vast majority of people on the planet, there is no economic mechanism to obtain access to that supply even though there might be perfectly valid reasons for that access (e.g., RTGs would remove the need to plug in your electric car).

Vest: In this particular context, the same incentives that are presumed to bring supply into the market will also assure the kind of quantities and prices that will prevent many prospective future Internet contributors from ever enjoying the same level of self-determination (translation: incentives to contribute) that has been enjoyed by every directly participating institution (a.k.a. future IP address dealer) to date.

Conrad: Actually, the fact that the IPv4 free pool is near exhaustion is what would (arguably) prevent the "many prospective future Internet contributors from ever enjoying" their access to porn ... err, contributing in a self-determinant way.

Vest: I think that misconstrues the situation. The specific condition that would indeed impose all of the "bad" effects, and that would result either from resource transfers or from doing nothing at all is the absence of a "sufficient" quantity of transparently interoperable IP (full-stop) addresses of any kind.

Conrad: You are treating the creation of a market as a disease when in fact it is a symptom. The creation of a market is a side effect of the continued demand for IPv4 addresses in the face of decreasing availability. The specific "bad" effects you are concerned about WILL occur regardless of the existence of the address market simply because the free pool of IPv4 addresses is being exhausted.

Vest: Actually I think you just restated what I said with different emphasis. In any case it sounds like we agree that transfer markets will do almost nothing to meliorate the problems that will inevitably accompany the disappearance of "sufficient" public/interoperable IP addresses of any kind.

Conrad: Yes and no. First, public/interoperable IP addresses are not disappearing. What will disappear is the availability of IPv4 addresses at a predictable cost. **The transfer market will encourage increased IPv4 address space utilization efficiency which will extend the useful lifetime of IPv4, albeit with significant change to the 'ecology' of Internet service provision.**

Vest: Here, again, it seems to me that we agree. I just go the extra step and infer what kind of changes in the ecology of Internet service provision are most likely.

Conrad: I also believe that RIR-mediated transfer markets will allow the RIRs to continue to provide a set of useful and valuable services to the Internet operations community in that they give ISPs a venue where they can (in effect) negotiate informal multi-lateral agreements regarding who is permitted to announce what. If the RIRs continued to try to play King Canute, all they would have succeeded in doing would have been to make themselves irrelevant.

Vest: I have tried, as much as possible given my circumstances, to completely discount the impact of any of this on the RIRs as such, and focus instead on the critical functions that the RIRs were chartered to perform or support (uniqueness, identification, routing capacity conservation, address resource stewardship). None of these functions will lose any of their significance as a result of the exhaustion of the unallocated IPv4 pool, but each will be much harder if not impossible for anyone to sustain on a self-policing basis because of the rejection of IPv6, and the unavoidable coordination

work that would be necessary to make IPv6 viable as a extension of IP addressing in general.

Conrad: However, in the long term, I agree that transfer markets can't alter the simple fact that there aren't enough IPv4 addresses to go around.

Vest: This is why I used the term "stewardship" in connection addressing above, rather than "conservation." Back in the early 1990s, when there wasn't enough IPv4 to go around in the form of classful prefixes, the community adopted a hybrid technical/administrative solution to relieve the shortfall. In the years after that, when it became clear that there still wouldn't be enough IPv4 to go around in the form of CIDR prefixes, DHCP and private addressing were introduced to manage the problem.

In all of these case the problem (well, one of them) was an address shortage, but the remedy involved technological means to absolutely reduce the "enough" threshold. Moreover, that "enough" refers usable public/interoperable IP addresses, not strictly speaking to IPv4. Had certain steps been taken, IPv6 could have been made viable at the margins, resulting in not only a gradual reversal of the long trend of "enough" compression, but also in the continuity of industry openness to future new entrants. Having dispensed with that option, we are fated to chase scarcity for the remaining useful lifetime of TCP/IP (i.e., the "bottomless and endless" market I referred to yesterday).

Earlier **Vest**: Transfers may delay the point when everyone everywhere is harmed absolutely by this fact, but only by introducing an extended period during which almost everybody is harmed absolutely, with the net difference being the possible emergence of a few severely hobbled new entrants, and the generation of vast windfalls for a few lucky incumbents.

Economic Consequences

Conrad: **I'm afraid I don't see the 'absolute harm' in organizations continuing to have the option of obtaining IPv4 addresses that are easily recognized as legitimately allocated, even if exercising that option costs (potentially significantly) more than has been the case in the past.**

Vest: To me this is like saying there would be no harm if the auto industry only made cars that got 20 mpg on average, and that there would never be harm no matter how expensive gas became -- and there would be no harm even if automakers had the technology to triple that mileage, but declined to do so because introducing that technology would cost them 5% of their current profit margin. The harm is systemic, the victim is everyone that depends on cars (or the Internet) as a factor contributing to overall economic growth and dynamism. The perpetrators may or may not be blameless, but the harm exists regardless.

COOK Report: How might this work? ATT gets a block allocation with 50,000 addresses routes half and SELLs the other half?

Vest: No, AT&T would have to be totally irrational to do that.

Conrad: Agreed. AT&T (and other service providers) are going to be buyers, not sellers. Sellers will be folks who have address space but aren't using it. In the 'traditional' RIR world, if you aren't using address space, you're supposed to return it to the allocating registry for subsequent reallocation according to justified need (i.e., "from each according to ability, to each according to need"). However, to date, this hasn't been all that successful in

recovering allocated-but-unused address space as there has been no real incentive for folks to return address space they aren't using. Now that the IPv4 free pool is being exhausted, the assumption is that the market will provide that incentive. If you believe recent studies, less than 10% of the IPv4 address space is actually being used.

Vest: Think about it this way: what if (any big service provider) had 50 lit 10Gbps waves along some fiber route today. Would you expect them to dump 25 at any price all at once just because they could? How do you think that would affect prevailing market prices and their own pricing power going forward? Here's the kind of thing that a rational IPv4 dealer would be more likely to do:

1. Offer 1 x /24 for $50,000 1a. If there is no interest at this price for a long time, offer 1x /24 for $45,000 and 1x /25 for $25,000... 2. If there is an enthusiastic response, offer the next 1 x /24 for 100,000 and 1 x /25 for $55,000 2a. If there is no interest at this price for a long time, offer 1x /24 for $90,000 and 1x /25 for $50,000... 3. If there is an enthusiastic response, offer the next 1 x /24 for 200,000 and 1 x /25 for $110,000 and 1 x /26 for $60,000... 2a. If there is no interest at this price for a long time, offer 1x /24 for $180,000 and 1x /25 for $100,000... 4. If there is an enthusiastic response, offer the next 1 x /24 for 400,000 and 1 x /25 for $220,000 and 1 x /26 for $120,000 and 1 x /27 for $70,000... Continue...

Eventually routability (or more proximately additional routing surcharges) will determine how fast deaggregation occurs, and prefix-level routing surcharges will be driven by (unpublished) inter-provider routing settlement arrangements, the baseline for which will be set by smaller group(s) of very large service providers.

COOK Report: Does the market selling process include a guarantee of routability for what you buy?

Vest: The "market price" for IPv4 transfers will likely include something the opposite -- a guarantee of uncertain routability.

Conrad: An implied guarantee of uncertain routability, sure.

Vest: However, incumbents will always be able to offer a guarantee of routability, called "being a dedicated customer."

Conrad: Not really. Same as today, they will only be able to guarantee routing within their own infrastructure. In particular, if you as a customer of a particular ISP present them with a random IPv4 /32 that you obtained off eBay for $50,000 (or whatever), your ISP may accept that /32 and may even announce that /32 to their peers, but there is no reason those peers (or their peers, if the route in propagated) will accept that announcement. I believe this fact will tend to put a damper on some of the wild excesses some folks have been concerned about.

Vest: **The fundamental break that IPv4 exhaustion/transfers represents will not affect routability or even reliable routability per se -- not immediately anyway. Rather, it will affect who is capable of being/becoming a credible provider of (mostly) reliable routing services for large numbers of customers. If you get IPv4 the old way, you may be able to join that club. If you are not lucky enough to be a member of that group, then you can still become a reliable customer of more-or-less reliable services provided by one of the members of that club.**

COOK Report: Or does ATT establish a new preferred

customer class and says for x hundred a month you get your own IP number which we will route for you for the next 12 months for Y dollars a month and at the end of THAT 12 months --depending on market justification -- the monthly price to keep it routed increases to Z dollars. So you have a money creation machine by which the most backward telcos can maintain their business models by squeezing the little guys out. Unless you grandfather existing folk like us, the day will come when Verizon tells its FiOS customers they have a MONTHLY ADDRESS surcharge? (By the way, I have looked at my IP address and although I thought the purpose of DHCP was to change the IP number on a regular basis, the IP number is remarkably stable.

Vest: No doubt price increases will percolate down to everyone, but in the retail customer examples you cite, the customer enjoys no control whatsoever over the IP address that ultimately mediates their attachment to their Internet access provider. **The big change - that will not directly affect them, except insofar as the non-existence of additional new public IP addresses - will mean that the chances of some completely new access provider, one that is not merely "reselling" services provided by the current incumbent, emerging to compete with their current provider get than much closer to zero.**

Conrad: To my knowledge, we aren't seeing a whole lot of those today when IPv4 address space is essentially free.

Vest: How about YahooBB -- the first or second-largest Internet access provider in Japan since 2001 or thereabouts. Didn't exist as such at all before 2000.

But more generally, for maybe a century leading up to 1980 there were about 200-300 "top-level" communications services providers in the world -- roughly one per nation-state. Today there are at 30,000+ if BGP is any guide. Not all are really autonomous, and some do nothing but provide an administrative overlay for someone else's facilities, and the majority provide connectivity only for themselves (i.e., they achieved "self-determination", at least in this narrow context), but even so there are probably at least an order of magnitude more than there would have been otherwise, largely because of TCP/IP and the technologies that only the demand created by TCP/IP (or some other similarly useful bypass technology) could have produced.

But as we all know, and I've said here and elsewhere many times, facilities or infrastructure are the locus of many enduring and completely unrelated bottleneck problems in many places around the world. **No amount of infrastructure can fix the disappearance of useful public IP addressing (unless maybe you can pull off a global monopoly), but for those facilities owners who might be interested in (re)establishing a more modest territorial monopoly, the disappearance of an effective bypass technology might not look like a such a bad thing at all. Better still is no particular effort is required to make it happen -- just do nothing different!**

COOK Report: How many Class A's does DOD have?

Conrad: At least a dozen.

COOK Report: Could a National public internet system have routing done in such a way that the address space were used far more effectively?

Conrad: In theory, sure. First, nationalize all the ISPs. Then, force everybody (ISPs and their customers) to renumber. Customers get to renumber into RFC 1918 address space and get assigned a single public IP address

from the national ISP that they NAT out of...

OK, maybe there is a flaw with this plan... :-)

Earlier **Vest**: Of course, at various points in history, changes in technology (e.g., CIDR, DHCP, NAT, et al.) changed our perception of what "sufficient" means -- **and no doubt most feel that technology has not kept up, with the resulting perception being that we haven't had "truly sufficient" address resources for a long time. But there have always been "some" for both incumbents and aspiring new entrants. No more.**

Conrad: Exactly. The IPv4 free pool will be consumed "soon". A market will encourage increased address utilization efficiency at some cost, defined by the market. Without a market, the cost will be infinite. There isn't really another option. I'd argue even the 'lack of market' option doesn't realistically exist -- a market exists now, the only question is whether or not the RIRs will continue to be relevant moving forward.

Vest: I also agree that the 'lack of market' option does not realistically exist either. However, the kind of market we've gotten ourselves into is the bottomless/forever kind.

Conrad: The market will exist as long as there is a demand for IPv4 addresses. I seriously doubt this will be forever.

Vest: Had steps been taken anytime before today to create the possibility of establishing full IPv4-IPv6 substitutability over time (maybe even a long long time), this might have been remembered as a unavoidable but temporary, self-healing break.

Conrad: You state this as if people intentionally decided to not establish full IPv4-IPv6 substitutability.

Vest: Not quite. Until today, [the RIPE address allocation policy] failure to establish full IPv4-IPv6 substitutability could be regarded as a passive, non-intentional choice -- no good word for this comes to mind, but the sort of thing that a lawyer or ethicist might describe as a something (crime, error, sin, etc. - an "outcome") of "omission." After today, the shift in incentives that resource transfers will engender will reinforce and convert this passive stance into an active, intentional, commercial strategy -- i.e., an outcome of "commission."

Earlier Vest: **As it is, address scarcity will now become the single most influential determinant of Internet technology, architecture, and market structure forever, until TCP/IP is replaced or bypassed by something else entirely.**

Conrad: I disagree. I believe the single most influential determinant will continue to be whether money can be made. Technology, architecture, and market structure will adjust.

Vest: Again, I think we said the same thing; I just used more colorful language ;-)

The Internet Under this New Policy Becomes the New PSTN

Vest: The Internet is the new PSTN :-(

Conrad: Yes, but I'm not sure how this is follows from your assertions.

Vest: I mean that a primary driver for Internet growth was the value of TCP/IP a mechanism enabling the bypass of lower-level bottlenecks. Leninists can start their victory parties now.

Conrad: The establishment of a market over "to each ac-

cording to need" is something Leninists would celebrate?

Vest: Sorry, this is a recurring misunderstanding of my own making. I am referring to Leninism as a strategy for political change, not the substance of Marxist or communist theory. Lenin's unique contribution was basically an early, mirror-image version of Naomi Klein's "Shock Doctrine" thesis. He fervently wished for/awaited the spontaneous collapse of economies/societies -- as a result of their own "internal contradictions" -- because following such collapses the people would be more receptive to his own vision of how things should be organized.

It's clear that some people (some list members included) are keen to see IPv6 fail not only because they find IPv6 itself distasteful, but more generally because they want to see TCP/IP abandoned, and absolute failure would provide fertile ground for a blank slate successor (or so they seem to believe).

Earlier Conrad: I actually think that fear is wildly overblown given everyone on the planet has over 18 million IPv4 addresses at their disposal for the cost of a NAT box and the availability of one public IPv4 address.

Vest: If that's really true, then maybe my concerns are completely unfounded, and PI IPv4 will be plentiful and cheap -

Conrad: Of course PI IPv4 will not be plentiful or cheap, but that's not what I said (actually the opposite).

Vest: Or maybe the market will start making individual /32s available at $50k each.

Conrad: Realistically speaking, can you think of a more effective way to get allocated-but-unused IPv4 address space back into play? Also, at $50K per IPv4 /32, IPv6 would likely become quite cost effective...

Vest: Cost effective for who, to do what?

Conrad: For organizations wanting Internet connectivity to absorb the cost of deploying new technologies (including NAT-PT or equivalent hacks to enable access to IPv4-only sites).

Vest: One could debate the ethics of this notion endlessly, but I'll gladly skip that step and point out that it won't work; the best thing it might possibly enable is a long death spiral. **Some organizations want general Internet connectivity -- i.e., to be consumers of general Internet-wide access. What they're going to get under the current circumstances is progressively more expensive, progressively more restricted, progressively more erratic/less useful Internet semi-access. Others might want to be producers of Internet connectivity and Internet access -- but unless they were chartered before today, they'll have to wait for whatever replaces TCP/IP sometime in the indefinite future.**

Earlier **Vest**: At $50k per /32, how much will it cost for a new entrant to become a viable (otherwise native IPv6-based) competitor for enterprise or consumer Internet access customers against the smallest /20-holding incumbent from the RIR era?

Conrad: How much will it cost the same new entrant if no market exists?

Vest: While completing that mental exercise, be sure to credit the incumbent with an additional $50k to commit to R&D, customer support, payroll, or stockholder dividends for each /32 that the new entrant had to purchase... **and be sure to remember that the IPv4 incumbent understands full well that his/her one sure unassailable market advantage is the high-confidence end-**

to-end reliability that only public IPv4 can provide...

Earlier Conrad: You appear to keep forgetting (or ignoring) that the only reason an IPv4 market exists is because of scarcity.

Vest: No I am not forgetting at all. What I am doing is making it harder for others to forget the fact that scarcity is being *chosen* in the presence of other alternatives.

Conrad: Sorry, what alternatives?

Vest: See

http://www.eyeconomics.com/backstage/GreenAllocationRegime.html

Conrad: The fact that RIPE-NCC has accepted a policy that enables a new form of transfer of address space does not preclude other policies.

Vest: Although this thread was occasioned by the passage of a RIPE policy, nothing I have written here is specific to the RIPE policy, or to any other specific RIR policy. **The *important* threat is not to the RIRs as they exist now, but rather to the critical functions that have enabled the Internet to grow so far/so fast and become so useful to so many. It's a contingent historical fact that those functions were originally installed in the RIRs, and there's a constant pragmatic debate about where they should be (or should have been) placed going forward.** My contention is that the path we have chosen is likely to make them absolutely unsustainable by anyone on the current self-regulatory model.

Conrad: However, I suspect all your "non-proposal" would have done would have been to set the ceiling on the cost of address space on the black market.

Vest: It could only have done that if it had worked. Most (not all) of the hypothetical advantages are logically entailed by each other.

Still might make for a good science fiction story though ;-)

Earlier **Vest** I am also less uncertain that you seem to be; this will certainly be the final death blow for IPv6.

Conrad: I doubt this as well. IPv6 will continue to be (essentially) free and the increased cost and uncertainty in obtaining IPv4 addresses will likely encourage additional IPv6 deployment.

Vest: The "increasing cost and uncertainty in obtaining IPv4 addresses" will never -- can never -- materially affect those that inherit IPv4 today from the RIR allocation era in the same way that it will affect everyone that comes after. The former will always enjoy choices about what kind(s) of addresses that they want to use to grow -- use NAT and private IPv4 (or NAT and IPv6) and end users everywhere pay a complexity premium, or use public IPv4 and the provider (and every other provider) pays a IPv4 transfer premium and a routing system cost premium. Non-incumbents won't be choosing between these alternatives; they'll be paying for both.

Goldstein: Transfer policy is likely to exist both officially and unofficially.

Vest: I think by this you mean that transfers will happen regardless of whatever (even voluntary) rules are (even voluntarily) adopted. I actually agree, so long as IPv4 is absolutely non-substitutable.

Among the "rules" that will be affected are the ones that have, to date, made IP number-related registration records useful for preserving uniqueness (i.e., avoiding unintentional address collisions), and facilitating inter-provider coordination, technical "event management,"

troubleshooting, and root cause analysis -- and on occasion, LEA. Maybe by the time the overall quality of those registration records falls so far that these functions are no longer supportable, address uniqueness will be irrelevant, and the Internet will have been purged of all possible risks of inter-provider technical failures and all possible benefits of inter-provider coordination, and LEAs won't even be interested in the Internet at all. Better hope so anyway, because nature abhors a vacuum.

Goldstein: If there's too much friction in the former ("we're making it hard because we want the world to go to v6"), then the latter will happen ("here, you can buy this subsidiary which we created just to own that /24"). There's an old Internet tradition about working around damage. Why should ICANN be viewed any differently?

Vest: Friction is a feature of all real vs. imaginary things, but otherwise you're right; the mere existence of any externally defined procedure will regarded by enough as too much bother to guarantee its overall failure over time.

And "working around damage" is not a tradition; it's a metaphor for how the TCP/IP overlay enables flexible bypass between interoperable end points (i.e., the kind that have public IP addresses) given any existing path between the endpoints. We're going to miss that reality, although the metaphor will no doubt continue to be used widely.

Conrad: ICANN has essentially no role here. The _only_ thing ICANN does (other than approve new RIRs which isn't relevant here) is approve global allocation policies that have already been approved by all 5 RIRs. There is no consensus (to put it mildly) among the 5 RIRs on this particular issue, thus ICANN is not involved.

Earlier **Vest** And in the absence of any technical factors limiting the proliferation of routing table entries, or any new source of publicly useful IP addressing, every important aspect of any/every conceivable future will be shaped primarily by some combination of oligopolistically imposed administrative rules and/or externally imposed laws/regulations.

Conrad: You forgot the "WE'RE DOOMED!! DOOOOOOOMED!!" bit... :-)

Vest: Hey, I volunteered that maybe it won't be so bad. And I've also repeated (often) that I would like nothing more than to be proven wrong. I'd like it even better if we didn't have to wait to find out the hard way, if someone stepped up right now with a persuasive, internally consistent argument that happier outcomes are more likely. I'm still waiting...

Goldstein: I'm not going to forecast happy or sad outcomes, though I have some faith that users and vendors can figure this out.

I am concerned with discourse that suggests that the only solution is IPv6, and that the only question is how to get there. IPv6 is a blind alley. The marching band is already piling up at the end, and beating the drums harder won't open it up.

Vest: I think you've made your opinion on this point clear.

Goldstein: And by the way, I hear that David Meyer has just noticed some problems with LISP and path discovery, showing that Loc/ID doesn't scale as expected, and thus neither does IPv6. Nothing that couldn't have been discovered ten years ago, but nobody was looking.

I'm positing that IPv6 simply won't do.

Conrad: Pragmatically speaking (given the "vast installed base" (to borrow a term :-)), for the foreseeable

future we have two alternatives:

- IPv6 + NAT-PT (or equivalent) - IPv4 + NAT

Of these two, I personally believe IPv6 has a slight edge if only because it is (essentially) freely available. Note, however, that they aren't mutually exclusive, and in fact, realistic deployment scenarios will involve both.

If the Alternatives Are Bad Enough – Build Something New

Goldstein: Thus the solution has to be found elsewhere. Why don't we actually start looking for one?

I suggest that this best be viewed as a competition of ideas, not a single committee process (like, uh, IPNG). Let a hundred flowers bloom, and let's see what the bright people out there come up with. My guess is that a better solution can be found which can be architected and implemented more rapidly than IPv6. (But then going straight through a blind alley always takes a very long time.)

Vest: That sounds a bit like what I asked for previously -- a plausible, internally consistent scenario that does not depend critically on wish fulfillment, magical thinking, or faith that everything always work out in order to arrive at some happy end state.

So, let's have one!

Cole: Could someone explain the size of the "bottleneck" problem? What if, say Chris Savage and Tom Vest formed a company that had ONE of the old-style, public IP numbers (IPv4), and the world's largest NAT -- using something like IPv8 or so, so that everything beyond that one IPv4 number had its own, unique number, up to "billions and billions" of devices. With unlimited hardware and software resources (obtained from renting those IPv8 numbers), what are the constraints on capacity through the Vest/Savage port to and from the Vest/Savage world?

I pick Vest/Savage so we can assume the operator of this "crucial" IPv4 port is COMPLETELY altruistic <grin>.

Vest: I make no claim to complete altruism, or to any altruism at all. I'm interested in what will enable the Internet to continue becoming more useful in more ways to more people. Unlike Fred and some others, I think that this is best/most readily accomplished by building on past achievements rather than by attempting to start with a completely blank slate and only match but exceed everything that could be possible through evolutionary change starting form here. I have doubts about the alternatives that are perfectly captured by something Harold Feld wrote in his latest blog entry:

"Be prepared to regulate where reasonable understanding of market incentives predicts that the most likely outcome is not what you want. Otherwise, take what comes from the Gods of the Marketplace, who remain both predictable and utterly indifferent to your needs or desires."

Just to prove my point about altruism, I'll answer your question as soon as you credit my Paypal account with $1000.

Will send account details under separate cover ;-)

Cole: If I had a $1000 to spare, I would get a computer that could play HD video (my 3-year old 3.0G Pentium turns them into slide shows). Then with my 2nd $1000, I would pay for your answer.....

Vest: Okay, maybe I'm no altruist, but I am a sucker for a PC hard-luck story ;-)

If your single /32-holding network enterprise is willing and able to find any peers/upstream providers to interconnect over IP UNNUMBERED links (which I believe

is basically unheard of today), then the prevailing assumption is that they would be able to use that single IPv4 address to "comfortably" provide simultaneous NATed service to a single in-house cluster up to 65 RFC 1918 or IPv6 client "desktops" under current conditions. If port utilization per user is extremely low by current standards, maybe this could be extended up to as many as 650 desktops -- and absent further NATing, in a year or two the max might conceivably be down to 30/300 or less.

Cole: OK -- 65 to 650, depending on how much the end-users use the system

Vest: If your single /32-holding network enterprise is interested and able to find any "pure IPv6" downstream transit customers who do no address translation of their own, and who are willing to pay for transit from a provider that uses IP UNNUMBERED to connect to the rest of the Internet, then you could substitute "enterprise customer" for "desktop" in the above description. But your hypothetical /32-based new entrant will competing for enterprise customers with incumbent IPv4-based service providers that will be able to easily provide one (maybe even more than one!) IPv4 to each of their direct customers, so this is never going to be a viable business for him.

Cole: -- but -- the supposition is that those already holding IPv4 numbers either have assigned all of them (we have "run out") or are demanding two children and a grandchild to be named later for each of them. If that is the case, I am thinking "crummy and cheap" might beat giving away descendants.....

A Market Will Come - Only Question is What Kind?

Vest: So far, what I've been describing is a scenario in which all future new entrants will be doing *both* -- i.e., giving away descendants (i.e., many $$$ for few public IPv4 addresses) in order to achieve just the possibility of offering crummy and cheap (i.e., delivery of once-translated packets to a small number of clients, or highly iffy (n)-translated packets to a larger number of clients).

If this sounds like a not-economically-viable scenario to you, then I have communicated my meaning successfully.

Alternately if Fred is right about both a future of cheap and abundant provider-independent IPv4, and the likelihood and benefits of full NAT standardization, then things will turn out differently.

Earlier **Vest**: There might be more potential multiplexing opportunities in there that I've failed to consider, but each one would result in a relative increase of complexity and loss of the kind of visibility required to quickly diagnose and remedy problems. At the level that your hypothetical single-IPv4 player will operate, the market will permit a variety of business models ranging from providing minimal service to a very small number of customers, to providing oversubscribed and increasingly dodgy minimal service to a somewhat larger number of customers.

And the current "full service" IPv4-based incumbents will be competing right alongside them.

Cole: Finally, there is a question of overlay -- if Verizon on occasion, and munis often, are willing to do physical overlay with fiber where copper already exists, could some other organization similarly start "overlaying" some at this point mythical new address system over the old. In the early days, end-users often had 2 phone books and 2 phones -- and when I had a nonprofit to buy such things for me, I had both DSL and Cable, while I was picking one. So I could see a transi-

tion being "ugly" at the edges -- this connection goes to subset A, this other one to subset B -- for a while, perhaps a long while, but constituting a transition to a new system nonetheless, even without some universal agreement to do so and over the strong opposition of the incumbents in the older system. (But then, I am the list's resident optimist....)

Vest: Anything is possible I guess. **Standardizing NAT would presumably permit successive translations with less iffy results. But of course provisioning that widely would almost certainly entail the same sort of global/ubiquitous hardware upgrades that IPv6 would require, and which have been rejected as too expensive/too much trouble by the service provider community.**

Conrad: The free pool of IPv4 addresses is being exhausted. The implication of this is that the cost of obtaining those addresses is going to go up. You appear to be suggesting that a market should be disallowed because the market would set the price.

Vest: Technically, my own humble non-proposal would not have "disallowed markets" to set prices. Right now (or at least before yesterday), there is no such market. More generally, there's no such thing as the pure, monolithic, canonical "Market" outside of the realm of mathematical abstractions. **The point of the debates of the past 18 months has been to consider *what kind(s) or market(s)* should set prices. The "green allocation regime" idea would have not have created or reaffirmed any restrictions on decentralized resource transfers. What it would have done is establish a mechanism, which could have coexisted alongside whatever else was going on, that would have assured a gradual elimination of the bottleneck requirement for IPv4 addresses, while at the same time preserving all manner of other functions that seem to me to be non-optional (no need to repeat them here again).** Even my uncanny powers of foresight don't provide an answer to whether that bottleneck would be relieved through the vehicle of IPv6 adoption, or something now unimaginable that would be even better -- but it would have been eliminated.

Conrad: If the market were disallowed the cost of obtaining address space that was registered with the RIRs would become effectively infinite.

Vest: Please stop subjecting me to this straw man criticism. The problems I have identified are endemic to the non-substitutability of IPv4. We may disagree about whether resource transfers will marginally aggravate or marginally meliorate these existential problems, and at what collateral cost, but **the bottom line *that we both agreed on already* is that "the 'lack of market' option does not realistically exist."**

Conrad: This seems to me to be more harmful than allowing people to meet the price the market defines since it means that everyone, including your mythical 'new entrants', would be unable to obtain registered IPv4 addresses despite the fact that up to 90% of the available IPv4 address space is unused.

Vest: **Under the arrangement I described, this harm would not exist for new entrants, because a small RIR reserve would remain available for "lifeboat" size allocations as long as necessary in order to get to whatever comes next. And since nothing would preclude incumbent IPv4 holders from trading amongst themselves as they saw fit (after all, nothing can be done to stop it, we have con-**

cluded), this claim has no basis.

Earlier on December 17 **Conrad**: I'm afraid I don't see the 'absolute harm' in organizations continuing to have the option of obtaining IPv4 addresses that are easily recognized as legitimately allocated, even if exercising that option costs (potentially significantly) more than has been the case in the past.

Vest: To me this is like saying there would be no harm if the auto industry only made cars that got 20mpg on average, and that there would never be harm no matter how expensive gas became -- and there would be no harm even if automakers had the technology to triple that mileage, but introducing that technology would cost them 5% of their current profit margin.

Conrad: Sorry, I don't see how your analogy applies to my comment. Let me try again:

The free pool of IPv4 addresses is being exhausted. The implication of this is that the cost of obtaining those addresses is going to go up. You appear to be suggesting that a market should be disallowed because the market would set the price.

Vest: After today, [Ripe's approval of a v4 address market] the shift in incentives that resource transfers will engender will reinforce and convert this passive stance into an active, intentional, commercial strategy -- i.e., a something (crime, error, sin, etc.) of "commission."

Conrad: So, just to be clear, you are saying that the existence of a recognized market in IPv4 addresses is going to terminate innovation that would allow folks to make more efficient use of IPv4 address pools, translate from IPv4 to IPv6 (or whatever), and provide other technologies that allow ISPs to continue to grow?

Vest: No. **What I am saying is that once IPv4 becomes a unique, non-substitutable commodity -- the market for which is completely locked up by current incumbents -- the balance of incentives that drive technology development will tip toward advances that preserve IPv4's artificial value as a bottleneck input, and away from advances that would tend to eliminate that value. Multiplexing advances like carrier-grade NAT are a perfect example -- they make ownership of IPv4 even more valuable. Once providers start bankrolling and deploying products like that in earnest, there's no way that they're going to want to reverse course and support steps that eliminate that bottleneck. And so we'll be stuck with a new non-bypassable protocol-level bottleneck forever --- or at least until TCP/IP is supplanted.**

Conrad: I actually expect the exact opposite.

Vest: Well, I don't understand how you could, realistically, but I do hope that you are right!

[**Editor**: Reverting to a thread with Rollie Cole and Tom Vest several hundreds of words above.]

Vest: If your single /32-holding network enterprise is willing and able to find any peers/upstream providers to interconnect over IP UNNUMBERED links (which I believe is basically unheard of today). . .

Conrad: Why would you need unnumbered? Pretty much any ISP is willing allocate a single /32 address. You don't need PI space for what Rollie is suggesting.

Vest: Pretty much any ISP is willing to give you a /30 for p2p for your peering link today. Maybe the exhaustion of IPv4 won't cause transit providers to reclassify that as a "premium feature" with its own recurring service fee--but maybe it will.

Conrad: Of course they will. However, this isn't relevant to what Rollie was asking, as you note.

Vest: In any case, Rollie asked for an example of what an independent service provider could do independently with *just one* /32.

Conrad: Perhaps the disconnect we're experiencing is due to assuming different models. I am assuming the /32 is the network attachment point joining the Internet and Rollie's "IPv8" network. This attachment point does not need to be PI. If you assume "IPv8" = "IPv4", this describes most of the networks connecting the Internet today (that is, SOHO networks connecting through a NAT box). This model of interconnection allows for O(2^56) devices (assuming a single layer of NAT). While this model breaks the "end to end" architecture and adds significant complexity if you wish to offer services on the "IPv8" network, it is a perfectly rational and workable model that has some interesting benefits (e.g., changing Internet providers is trivial). It isn't the network architecture I would choose, but then again, I bought Beta video recorders.

How expensive will that single /32 be in the face of competition amongst multiple IPv4-only Internet service providers? I am supremely skeptical it would be $50,000. In fact, I'm skeptical it wouldn't be bundled with the connectivity the ISP provides as is done today. However you appear to be assuming something else.

Earlier **Vest**: then the prevailing assumption is that they would be able to use that single IPv4 address to "comfortably" provide simultaneous NATed service to a single in-house cluster up to 65 RFC 1918 or IPv6 client "desktops" under current conditions. If port utilization per user is extremely low by current standards, maybe this could be extended up to as many as 650 desktops -- and absent further NATing, in a year or two the max might conceivably be down to 30/300 or less.

Conrad: Where are you getting your numbers?

Vest: IPv4-IPv6 coexistence materials, mostly from the IETF/IAB. Stuff like this: http://www.nttv6.jp/~miyakawa/IETF72/IETF-IAB-TECH-PLENARY-NTT-miyakawa.pdf

I've seen some current estimates suggesting that an individual user (e.g., someone using Google apps and similar things), could generate peak concurrent demand for as many as 1,000 TCP ports. 1000 ports per user conservative, 100 ports for grannies or with contention -- 65 to 650 simultaneous users.

Conrad: I make the assumption that web (and other) services providers, at least the commercial ones, will choose to implement their services in ways that allow them to reach as many paying customers as possible. As the model of connectivity to the Internet evolves, I believe services providers will take into account that if they open up 1000 ports, their customers who are behind low performing NATs (or whatever) will get crappy services and choose a competitor that provides services that work better with their low performing NATs (or whatever). Maybe I'm wrong and folks like Google will write off the folks who are behind NATs.

I guess I assume that people and companies will adapt to the interconnectivity model that the Internet evolves to as they have done in the past (e.g., 20 years ago, the idea of NAT on the Internet would have been laughable). While I anticipate there is going to be a period of turbulence as everyone settles around a new strange attractor (so to speak), I do not believe the acceptance of a white market in IPv4 addresses signals a capitulation to the "address space oligopolists" as you would seem to imply.

I suppose we'll see in something like 796 days.

Earlier **Conrad**: You are aware, of course, that there are entire countries where, due to regulatory stupidity, normal consumer-level access is provided through multi-layer NAT, right?

Vest: I am abundantly aware, just as I am aware that there are entire countries where "independent" local ISPs cannot obtain more than a /24 of "provider dependent" IPv4 from their upstream provider, which makes it impossible either to change providers or multi-home. In the former case, even if consumers never ever see a public IP address that doesn't mean that the beneficiary of that regulatory stupidity (or the more remote beneficiaries -- the offshore provider who services the whole thing as a "service export") has a substantial block of IPv4 dedicated to mediating those connections.

In any case, if this is your existence proof that multi-layer NAT can work, I guess that means that we should also expect the same kind of kludgy, half-crippled service to become the norm in our -- everyone's -- future? I couldn't have made the case better myself ;-)

Conrad: You are aware, of course, that router vendors are busily testing "Carrier Grade NAT", that can handle _vastly_ more than '65 RFC 1918 or IPv6 client "desktops"' right? Even today, you can buy off the shelf boxes that handle tens of thousands of simultaneous NAT sessions easily.

Vest: Hopefully the above clarifies.... I guess the crux is credible estimates of concurrent TCP port requirements... pls. let me know if you've seen dramatically different ones...

Rollie Cole Summarizes

Cole: Thank you all -- my question has provoked more interesting discussion than I could have hoped.

Let me offer a few predictions, mainly to see if I understand.

1. Nothing much happens, as Tom Vest suggests, until we are "out" of free IPv4 addresses (the historical pattern of waiting for rain rather than fixing the roof in advance).

Conrad: Not quite. I believe one advantage of the RIPE policy change is that it sets the stage for a change in the 'ecology' prior to run out.

Cole: 2. Once out, three things start to happen, at varying rates of speed and with various degrees of success:

2.1 a "market" develops in IPv4 addresses, with them that has gets -- the more you have, the more you can both buy and sell

Conrad: A market already exists, albeit black or at least grey. The RIPE policy announcement that triggered this discussion is seen by many as the first step in turning that grey/black market white.

Whether or not you have address space does not determine whether you can buy address space. However, pragmatically speaking, it will be the ISPs that will have the most incentive to obtain as much address space (in as large contiguous blocks) as they can.

Cole: 2.2 technology to "extend" the use of those you already have (multi-level NAT et al) now gets really serious, even in areas where it was not favored before

Conrad: Folks are already quite serious about multi-level NAT, CGN, etc. Some folks appear to even be increasingly serious about IPv6.

Cole: 2.3 technology to "get around" the limit (IPv6, something Tom might like

<grin>, et al.) begin to receive serious attention, but may or may not begin to be used

3. the whole system is much more chaotic for a while, perhaps a very long while, until and unless a less chaotic system begins to gain traction

Conrad: Perhaps it might be that there is a new axis of chaos in addition to the chaos created by governmental policy and economic drivers.

Cole: 4. **Although collectively we might have the technology "smarts" to reduce the time spent in steps 2 and 3, we lack the institutional/social "wisdom" to employ it. (This appears to me to be the crux of Tom's point, but I am happy to be corrected.)**

Conrad: I'd agree in the sense that there has been essentially no incentive to apply the technology 'smarts'. Where Tom appears to believe the implementation of the RIPE policy will stifle this application, I actually see the policy as necessary to incentivize innovation prior to hitting the brick wall of run out. I believe you're already seeing the application of technology 'smarts' in industry efforts like CGN and the renewed efforts at translation between IPv4 and IPv6 in the IETF as folks realize the brick wall is approaching.

Vest - **To Rollie Cole:: Very well summarized.**

The only thing missing are the interconnections -- i.e., 2.1 leads to 2.2 lock-in, and as 2.2 continues, 2.3 becomes increasingly impossible -- at least any version based on any form of TCP/IP that is currently available or under development (AFAIK). So what might look like a natural, evolutionary process is really a giant leap of faith -- not into the unknown, but rather into a known bear trap, but with faith that it won't bite, this time, maybe...

Cole: Remember the principle of "outrageous profits." Change is hard, so does NOT happen until either (a) current situation is really, really, really bad; or (b) perception that new situation is enough better to justify the costs of transition. It strikes me that IPv6 fails on both counts at the moment -- we have not hit the ground yet in our fall, and IPv6 is not "enough" better. But both the awfulness of the current situation, and the perception of how much better something new might be can both change quickly.

Vest: There's no denying the fact that your description is consistent with the observable reality. And I know you are making an statement, rather than stating a preference or aspiration -- but I can't help but note how closely this understanding of change resembles the Naomi Klein "Shock Doctrine" hypothesis, as well as what I previously called "Leninism." As you rightly note, once things get "really, really, really bad," people are often willing to consider formerly unthinkable options. But not all of those options are certain to lead to "better." There are likely to be lots of "revolutionaries" with mutually incompatible visions vying for influence then, just there are now - just as there always have been whenever times get tough.

And in the mean time, things are still just "really, really, really bad..."

A Different IPv6 Alternative

On December 21 **Hendrik Rood**: I think your sketch is one plausible scenario, but consider another.

In pockets of the communication industry with clear foresight on the need to deploy vast address spaces an internal use of IPv6 starts.

This usage in pockets nicely fits in with commercial walled garden strategies and a de-

sire to avoid full exposure to the public Internet for certain information and communication services.

Deployment of IPv6 based services is driven further due to operational cost bases that like to internally exploit IP technologies for its low cost due to its large scale economies, while using the vast address space to avoid "collisions in address-space" that might hamper future merger operations or sale of country divisions.

Gradually a world of IPv6 deployment appears that is screened at least from the public Internet by application gateways, IPv4 services then become tunneled over the IPv6 infrastructure, consumer end users are not offered direct connectivity on IPv6.

Gradually businesses are offered IPv6 connectivity into the operators networks, but at steep tarif mark-ups and under strict conditions with respect of what applications can be run over IPv6, all in the name of isolating the current IPv4 insecurities.

Tier-1 backbones gradually expand their IPv6 routing abilities, but only on a transit base, too selected customers and far a nice markup too.

After a number of years people start to grasp that most of the cash flow due to servicing business networking and monetisable communications services is not over the public IPv4 Internet, but IPv6 networks that are effectively a complex set of isolated networks and clusters of networks all with various kinds of far more elaborate bilateral compensation schemes than todays bilateral peering and transit agreements.

IPv4 remains the "experimental" network and the platform for mass distribution of content and services over 'unreliable networks' toward consumers and SOHO.

I think this is a quite plausible scenario too. Comcast deploys IPv6 internally, mobile operators deploy it, the military would prefer it that way too as it reduces their networks exposure for some threats, academia is due to it's own acceptable usage restrictions a specific market segment and ultimately Tier-1 backbone operators are finally able to get to offer a clear differentiated service. Some nation states may like the idea to become entirely IPv6 inside their country but operate a big IPv6-to-4 gateway and screening system a.k.a. big national firewalls.

Due to shrinking IPv4 address space, operators start to stear away from IPv4 and let the big address block trading there begin, as it educates end users very well to pay more and more for Internet access. There own new developments will all be IPv6 based, but walled.

Regulators will be very happy with this new situation too as the quite professional IPv6 space creates an area that can be easily 'characterised' as wholesale and business-to-business with a rather limited group of market players that is relatively easy to regulate and a vast new demand for their public services of regulatory interventions.

The net result is not that IPv6 will not take off, however it will be used as a tool by operators/large ISPs/governments/regulators to get rid of many of the aspects of the current openness of IPv4 that allows for easy bypasses, high level of competition, which depresses demand for their main services.

IPv6 will not be easily accessible by end users except as a premium service.

On 12/22 **Cole**: I agree that your scenario is plausible as well. A countervailing issue, however, is what I see as the growing "virtuality" of the large organizations that would be the operators of these closed IPv6 spaces. If a large organization wants to outsource to get cheaper la-

bor, or labor unavailable inside the organization, it is going to have to have a form of "openness" that will tend to press against, however strongly, a closed system. Maybe some organizations (and some countries) are willing to keep closed or pay a premium to jump through hoops for outside input, but many might simply open their networks a bit.

This tension between open and closed occurs in many dimensions -- including actual immigration of people, importation of goods, etc. I suspect it will be operating in the arch-econ of networks as well.

Vest: Degree and *direction* of openness versus closedness will always be a voluntary and tunable feature for "incumbents," i.e., those that have (relatively) abundant IPv4 by way of the RIR-administered, needs-based allocation system. That is to say, "incumbents" will always enjoy the option to implement whatever flavor of "virtuality" that they like with whomever they wish, provided the second party is willing.

At the very least, that freedom will not be a given for any who come after the era of needs-based allocation has passed.

Cole: I read a lot of Westerns, primarily for the descriptions of a geography I love. But the issues surrounding what happens when a lot of newly-arrived farmers want a piece of the big ranges "owned" by the ranchers who took it away from the Indians and/or Mexicans is a recurrent theme. Eventually, with or without a great deal of violence, the farmers usually win -- but not always. The IPv4 incumbents seem awfully similar to the ranchers in this plot..

Vest: .. and with that we are back to the U.S. west vs. Brazil vs. Rhodesia/Zimbabwe... or perhaps the Enclosure Acts vs. Shock Therapy in Chile vs. Shock Therapy in Russia...

Not to worry, I'm sure whatever happens will be regarded as rational, or at least inevitable, with a couple of decades of hindsight buffering.

Cole: That is the most likely scenario, but not the only one possible, and certainly not the most desireable. Here's hoping "evangelists" like Tom Vest can help us improve the transition at least a bit over historical precedent.

And So Does Harold Feld (Independent and Parallel Summary)

Finally **Harold Feld** (a List member) on his blog wrote a provocative summary: http://www.wetmachine.com/item/1428

RIPE Makes Me Vaguely Uneasy By Creating Legal Market For IP Addresses.

Talk to anyone who was involved back in ye olde days of the development of the Internet address system and underlying protocols and they will tell you that most of the major stuff - like the division of the domain name system into generic Top Level Domains (gTLDs) and country code top level domains (ccTLDs) just evolved on their own. Sometimes this worked out real well. Sometimes, not so much. But for better or for worse, these decisions set the pattern for how the internet evolved and created huge policy issues as the internet scaled up from a universe in which everyone knew everyone else to a system of global communications that always seems to be lurching toward - but never quite reaching - total collapse.

I'm not saying I could do better, or that anyone could. Indeed, I can argue that a lot of good stuff happened when people handled problems in

an ad hoc manner and that the major effort to put a little forethought and adult supervision over the whole process, the Internet Corporation for Assigning Names and Numbers (ICANN), turned into a total mess.

Nevertheless, it gave me a bad turn to read that RIPE-NCC, which allocates the IP addresses for the European Union, will now allow holders of IPv4 addresses to openly buy and sell these address allocations (you can read the policies around the address allocation here).

Why does this make me uneasy, especially when a gray market in IPv4 addresses already exists? Because it makes fundamental changes in an underlying piece of critical infrastructure. That always makes me queasy, especially when I know that those making the changes have not adequately considered the very many ways this can go badly, as well as the ways in which it can go well. OTOH, I also recognize that, as Ecclesiastes warns, "to the making of many books there is no end, and much study is a weariness of flesh." (Eccl: 12:12) Somebody needs to act sometime. Nor do I have a very clear idea what I would do instead to solve the IPv4 address exhaustion issue. **But I really worry about creating a class of powerful incumbents invested in preserving the value of their IPv4 real estate and opposing transition to IPv6.**

For more detail on this than any sane person would otherwise want, see below . . . .

I suppose my real problem is **that I just haven't dug into this area enough to really have an opinion. But then again, so few people have - which is part of what makes me uneasy. Few things rival IP address allocation in both importance and breathtaking, mind-numbing technical snooze-inducing power.** This makes it either the ideal laboratory of exciting new ideas or a veritable Devil's playground of possible mischief. Worse, I feel guilty for doing what I hate when someone else does it to me - swooping in suddenly after years of debate to try to go back to the beginning. OTOH, that's policy for you and I always have to suck it up. So, with huge apologies to Milton Mueller, here we go.

Lengthy But Necessary Background

Some background for those who have no experience with the issues around IP address allocation (lucky you). We human beings generally use domain names to find websites, or have other means whereby we let the technology do the driving. But routing actual packets of information relies on the Internet Protocol (IP) address. It is this long string of numbers that routers recognize as the actual address. Critical to the question at issue here, IP addresses do not need to have an associated domain name. Many, many machines need IP addresses to communicate with each other, and do not bother with an actual domain name.

Many years ago, the folks running the Internet (to the extent anybody ran it) adopted IPv4 as the standard for IP addresses. This created a world with a large number of IP addresses. Enough that the few people actually involved in these decisions at the time thought they would never run out. They had a delightfully informal system where a few people at the university of Southern California operating as the Internet Assigned Numbers Authority (IANA) just gave this stuff out to anyone who asked. After all, who cared? It was a couple of thousand people playing with enough address space to give an IP address to everyone in the entire world.

As the Internet scaled up to become a global communication system, some formality and structure inevitably came

into the system. Actually, the engineers running this, determined not to let governments get involved and screw things up, worked hard to create a system that was sufficiently international, open to interested parties, and insanely boring to most folks to keep administration of IP addresses pretty much a private matter. So today, we have 5 regional address registries (RIRs), one for the North American Region (ARIN), the South and Central American Region (LACNIC), the African Region (AFRIC), the Asia Pacific Region (APNIC) and the European Region (RIPE NCC). The IANA delegated address blocks to the RIRs to allocate to their regions, all under the loose umbrella of ICANN.

As long as everyone could get reasonable access to IP addresses at affordable cost, address allocation did not raise much concern. But as more and more machines needed IP addresses, the harder it became to just give addresses away. Other problems also started to emerge. The size of routing tables in the routers (the tables of where to send packets to reach their destinations) and the amount of capacity taken up by traffic information kept growing, and network engineers kept struggling heroically to keep up. Depending on how you look at it, the continuing functioning of the internet as a global medium under these conditions is either the ultimate triumph of Libertarian theory (because it is happening without centralized planning or government coordination) (except when it does) or a perpetual crisis waiting for the day of collapse when the ad hoc system of patches and solutions can no longer keep up.

So What's the Problem?

Which brings us to the present day. Despite clever and creative things, such as network address translation (NAT), to stretch the existing batch of IP addresses, the availability of IPv4 addresses is becoming increasingly more urgent. As we add more devices and more networks worldwide, we need more addresses. **If you can't get an address yourself, you need to hook into someone else's network and have them give you some of their IP address space. That works to some degree, but has problems for its own. If I want to offer a competing network to the large carriers, I may not want to get IP address space from them - and they may not want to provide it to me.** Even without the competition issues, needing to get IP address space from someone else rather than having a direct delegation from an RIR raises issues of cost and may also provide a way for the IP address provider to dictate what I can do with my network.

As a helpful example, consider the parallel addressing system for text messaging short codes here in the U.S. If I am a business wanting to transmit text messages to many people simultaneously and receive many text messages simultaneously - for example, to announce my choice of Vice Presidential candidate - I need a short code. I can only get one from the Common Short Code Authority (CSCA). To get this code, I have to fill out a fairly long form detailing why I want the short code, what I plan to do with it, and other rather intrusive questions about my possible business plan. If the CSCA doesn't like my answers, it won't give me a short code. If it gives me a short code, I have to obey all sorts of rules and guidelines. CSCA can inspect my business at any time to ensure I'm complying with the rules, and take away my short code if it thinks I'm not following the rules. Even when the system works, and no one leverages it for anticompetitive advantages, it is a real pain imposing significant costs on businesses trying to use short codes.

So far, IPv4 has avoided that fate. The RIRs generally charge modest fees for ad-

dress blocks to get folks to limit their use of address space to just what they need. Also critically, until now, they have generally prohibited the sale and purchase IP addresses from one private party to another. **The RIRs generally treat IP addresses like phone numbers, a network resource nobody owns and that must be available to everyone on an affordable, non-discriminatory basis. This does not prevent sub-delegation by carriers to customers.** In fact, the RIRs encourage this to save IP address space. Nor has it prevented a modest "grey market" in IP addresses from emerging. But the availability of IP addresses on reasonable terms has kept anyone from trying to arbitrage the IP address market or leverage IP address space for anticompetitive purposes in a major way.

But the continued depletion of IPv4 address blocks changes this dynamic. Despite everything engineers have done to stretch the IPv4 address space, the RIRs are running out. But the number of networks needing IP addresses keeps growing. We have now moved from a world where no one could ever imagine running out of address space to a world where we ask: how soon until we run out?

So Whatever Happened to IPv6?

As far back as 1998, the IETF adopted an initial standard for a new internet addressing system, IPv6, which would have even MORE address space. IPv6 has enough address space, I am told, for just about every molecule on Earth. IANA has delegated IPv6 blocks to all the RIRs, and the technical community has flogged converting from IPv4 to IPv6 for years. So why do we still have a crisis?

The biggest problem with the transition is that - despite significant efforts to address this problem - ***IPv6 is not backward compatible***. For those who do not immediately grasp why I put that in bold italic, it means that if you use an IPv6 network, you can only talk to other users of IPv6 and cannot communicate with users of IPv4 (at least not without much work). This is a classic [link]"network effect" problem. Why would I spend money to build an IPv6 network when everyone else I want to talk to is on the IPv4 network? The failure of IPv6 migration to date pretty much answers that question: "no reason, so I won't do it."

Again, an example from another area of US media and telecom policy illustrates the point. When Congress started the transition to digital television as part of the 1996 Telecommunications Act, it initially made the transition from analog television to digital television voluntary. It would only happen when 85% of stations started doing digital transmission. But no broadcasters wanted to invest in digital television program development or transition until people actually owned televisions that could receive digital signals. But nobody wanted to buy very expensive televisions to receive digital signals unless they needed to in order to get programming, which didn't exist because no one could receive the programming. Finally, Congress broke the logjam in 2005 (so they could hold the 700 MHz auction and raise big bucks) and ordered all broadcast stations to stop analog transition and move to digital on February 17, 2009.

We don't have that equivalent for IPv6. We have no one forcing network operators to move from IPv4 or otherwise make their networks IPv6 compatible. Nor is it clear how we could achieve that on a global - or even regional or national - basis. And believe me, a lot of very smart, highly motivated folks have spent a good deal of time trying to figure out how to make that happen. So, for the

foreseeable future, we will need to figure out how to manage the IPv4 space in a world of scarcity rather than a world of abundance.

Enough Background! What Did RIPE Actually Do? And Why Does It Make You Uneasy?

RIPE decided to adopt a private market solution. For the first time, if you hold an IP address, you can sell it to someone else, at whatever price you want, under whatever terms you want.

This may seem a funny time to try private market approaches, and Europe seems an odd place to engage in the experiment. But this goes to one of the key points about public policy - it gets shaped by those who show up to play. Remember how I mentioned how insanely complicated and boring this stuff is to most ordinary folks? That means a very small set of actors actually shows up to do the work, and they therefore make the rules.

So who shows up? Well, engineers of course, many of whom work for companies that care a great deal about address allocation. A private market solution makes their lives easier, as they already have IP addresses and would prefer to get them by just buying them from people who want to sell them. Yes, there is a long tradition that engineers check their company allegiance at the door when they show up for technical coordination. But to the extent that's true, the very experiences that make their insights useful also shape their results. These guys have a problem they want to solve, this looks like a fairly straightforward way to do it. And, no offense to engineers, public policy is as much a specialty as engineering. I may pick up some basic concepts, but I do not delude myself that I can run a network on my own.

A handful of public policy types do show up at these things, and tremendous credit goes to those who, like Milton Mueller have stayed involved despite the high cost, low return, and occasional outright contempt displayed in this community for disciplines like economics and public policy. However, as a combination of both law of averages and overall temperament, these public policy types are overwhelmingly free market enthusiasts who believe that private markets inevitably allocate resources better and more efficiently than any other system.

O.K., So What's Wrong With A Free Market Solution?

Actually, for all I know, this may prove the best way to handle the problem. As I have often said, there is nothing intrinsically good or bad about market mechanisms. Markets do work pretty well for distributing lots and lots of goods. Despite growing scarcity of IPv4 numbers, there are still an awful lot of them. In theory, that makes it very hard to corner the market, or leverage the scarcity through a cartel or even conscious parallelism. There is already a grey market in IP addresses, so trying to stop all trade in IP numbers is probably a losing proposition. Nor do I think that a public interest allocation system a la FCC license allocation could be easily implemented, or even necessarily a good idea.

Playing dice with critical infrastructure is almost always a losing bet: History shows that allocation of scarce critical resources through market forces almost always turns out to be a phenomenally bad idea. For one thing, despite all claims to the contrary, I have never met a critical infrastructure resource scarce enough to need market distribution but so plentiful it can't get leveraged at some point in the distribution chain. We got into this mess by assuming we had all the address space we could ever use. Now we assume we have so much the market will

never grow distorted. I see no reason why this assumption will prove more accurate than the previous assumption.

Keep in mind we do not need some single entity - or even a small cartel of evil greedy capitalists in a smoke filled room - to create real problems in market structure for critical resources. But to the best of my knowledge, no one has made a serious examination of the actual market or its implications throughout the supply chain. The closest is this analysis of IPv6 issues by OECD. [**Editor**: see http://www.wetmachine.com/item/www.oecd.org/dataoecd/7/1/40605942.pdf.] We're flying blind here, with the underlying address system of the Internet. Lacking a religious faith in the Gods of the Marketplace, this induces queasiness rather than the warm fuzzy that comes from embracing the Market in its Coasian Perfection.

It also really worries me that the entities best placed to take advantage of the new market structure are those with the greatest financial interest in figuring out how to leverage the market - the carriers. The ability to absorb enough of the market to force new entrants to accept disadvantageous terms for access to IP addresses worries me. Again, this would not require that carriers control all IP addresses. They just have to control enough to make it more expensive than submitting to terms. Anyone who studies how Microsoft gained its dominance in the software market should recognize that the ability to create just the right set of incentives and disincentives in provision of a critical input is key to establishing and maintaining market power.

Finally, the possession of critical infrastructure assets invariably creates opportunities for arbitrage. If there is money to be made by playing games, folks will work on figuring out the games to play. Enron built a fortune (in the short term) figuring out how to manipulate the private electric markets. Ask the people of California - or even the Enron stockholders - whether creating opportunities for arbitrage by adopting free market mechanisms served them well.

What happens to IPv6 Migration: We also have the peculiar paradox that the more we make it possible to stretch IPv4 address space to avoid a crisis, the more difficult we make it to transition to IPv6. Every dollar of invested cost in maintaining the IPv4 space by a company or new entrant is one more dollar of lock-in for the existing network. **Worse, anyone that controls IPv4 address space that has value has an incentive to maximize the value of that resource. As this resource becomes utterly worthless once a tipping point's worth of folks switch to IPv6, it creates a real incentive to delay the IPv6 transition as much as possible.**

Again, I come back to the fact that those best positioned to become this class of IPv6 resistant incumbents are the carriers. That's not a matter of evil intent. It is a straightforward economic reality of the fact that it is the carriers who do the routing.

Mind, if IPv4 incumbents became too powerful or obnoxious, it might drive a critical mass of people to adopt IPv6 to escape. The problem with this theory is that incumbents tend to be fairly good at finding a middle ground between "the competitive market price" and "so awful that people will pay anything to escape." We call this middle ground "extraction of monopoly rents." While nobody likes this very much (except those that have the critical asset), it costs too much (in both monetary and non-monetary terms) to do without and develop and alternative. Worse, during the period of shift, the incumbents can modify their extraction of monopoly rents to make things more livable

and reduce the incentive to develop an alternative to a point where it doesn't happen.

I do know that the longer the transition from IPv4 to IPv6 goes on, the harder it will be to do. Creating a class of actors invested in IPv4 as a valuable asset rather than simply as a question of cost avoidance will not help.

Does what happens in Europe stay in Europe: I'm glad only one RIR has decided to take the plunge. But how long will the other RIRs wait before opting for this approach? We had credit default options and subprime loans for many years before the collapse. Countries and institutions that initially rejected these instruments and investment vehicles as too risky succumbed to the constant pressure from investors eager to "get in on the action." Until last year, countries with conservative investment laws looked like fuddy-duddy dinosaurs preventing their citizens from getting in on the good times. This year, not so much.

I am worried that we will see a huge rush to declare RIPE's policy shift a "success" before we even know what success would mean here. I expect that "success" will be measured in something really stupid and irrelevant, like total number of transactions as compared to assignments. I also anticipate that as parties begin to make money from the sale of address blocks, similar parties in other regions will begin to lobby for these changes.

Without having some framework to figure out what the Hell would constitute "success" or "failure" of RIPE's market experiment, we will almost certainly see it become the dominant model. Why? Because it is the only model, and its supporters will loudly proclaim its success. With no metrics to define success, this will look very convincing.

So What's Your Bright Idea Oh Wise One?

I don't have one. Heck, for all I know, creating an unregulated market in IPv4 addresses with no possibility of regulatory oversight may be the right solution. It would certainly make it very difficult for any single government to unilaterally regulate the IP address market. It also addresses a real problem of address exhaustion that we face with increasing urgency.

But I have a very bad feeling about this, and I have come to trust that instinct. At least, I trust it enough to explain at great length why this makes me uneasy.

Editor's Concluding Summary

Unless one can think at high levels in terms of complex systems, one will never understand the cause for concern. Routable v4 addresses are the currency of the internet. Without them one cannot play.

As Harold Feld recognizes at the end **this critical resource is being PRIVATIZED with complete lack of oversight and of transparency.** The ultimate libertarian free market. Remember the RIRs are membership organizations whose staff do what the ISP incumbent members tell them. They have existed to govern assignment of numbers - we are looking at a situation where when the numbers are all assigned they have not much reason to exist and where in the meantime all sorts of opaque products could be created by the ISPs that belong to RIPE and soon by the ARIN and APNIC members

I am not expecting IP equivalents of credit default swaps. But one wonders if someone might not get the idea of selling routing insurance. Remember that these are units

of currency for the internet at least and in a closed non transparent market the temptation will be there for the actors to say their IP numbers are better than someone elses' because they have deal with "x" and "y" where the IP numbers held by "z" get priority. In other words the same sort of factors that laid the financial sector low (i.e., huge risks and profit opportunities accompanied by total obscurity/opacity to outsiders) also has the potential to wreck the Internet services industry.

Think about the mechanisms of CONTROL that a private market of IPv4 routable numbers puts in the hands of the incumbents. The primary significance of the private market is that it gives incumbents the absolute power to control entry in the market; you cannot play unless you secure their permission, in the form of the precious IPv4 inputs that only they can provide. The Internet is now telco-ized. Absent regulatory interventions (which don't exist in most markets), new entry is impossible unless the incumbent telco grants you permission, usually by taking you on as a subcontractor or junior partner as it lets you use its IP numbers..

Open access networks.... sorry no numbers left. Eventually we may hope the government will nationalize the system. But can any transnational system be agreed on? Doubtful. What is clear *now* is that migration to v6 won't happen in time to make a difference and that industrial strength NAT can turn into a new carrier controlled bottle neck.

There is no agreement on those involved but its ironic that the locus of the counter revolution is in Amsterdam. Maybe Gigaport can get the world on circuit switched lightwaves in time?

International Bandwidth, Pricing and the FCC Deliberately Suppressing Demand?

Hendrick Rood: When I went searching for some international bandwidth figures I bounced into this nice presentation by *Telegeography*

http://www.ptc.org/ptc08/participants/speakers/papers/SchoonoverFinalSlides.pdf

Alongside this Executive Summary it provides a nice overview

http://www.telegeography.com/products/gb/pdf/Executive_Summary.pdf

But it also gives you one pause:

Why is actual Purchased Capacity on lit cable systems (see slide 23 for Transatlantic) more than twice the amount of Used Capacity? It does not make economic sense in a market rife with overcapacity and with falling lease prices for 10G wavelengths to purchase capacity two years in advance and leave it unused.

The answer is, Purchased Capacity on lit cable systems isn't more than twice Used Capacity, but Used Capacity, in particular the International Private Line part is carefully underreported.

The cause of this is the FCC, who thinks it is smart to levy a regulatory cost recovery fee based on a LINEAR count of 64 kbit/s circuit equivalents in use.

See http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-284863A4.pdf and

http://edocket.access.gpo.gov/2008/E8-19899.htm

In particular attention can be paid to the filing's notes in the *Federal Register*: \45\ Petition at 7-8. Level 3 contends that this fee timing issue can make owners base their capacity turn-up decisions on non-market factors, such as activating circuits only at certain times of the year. Level 3 RM-11312 Comments at 5. \46\ Flag RM-11312 Comments at 6. Reliance observes that, with respect to high-capacity leases, the per 64 kbps circuit fee distorts the market. Reliance Reply Comments at 5.

Points are rather subtly stated but to grasp the results of the fee regulation discussed in the docket consider:

1. IPL users of big capacity links (e.g. 10 Gbit/s) pay more in regulatory fees than for the actual capacity leases
2. To avoid this excessive tax they shut down links each year in the last days of December, firing them up again in January, as on 31 December of each year accountants take stock for reporting.
3. The FCC pricing structure exempts common carriers, as well as IRUs and leases to common carrier use on private submarine cables.
4. It hits on corporate end users, academic research networks, IP backbones, independent ISPs, ASPs and content providers who lack common carrier status but need to operate global networks
5. Despite the ongoing discussion for a few years, the FCC postponed decisions to FY 2009 and will thus keep the current system in place for 2008.
6. As a result the Annual (Internet) Capacity Holiday Close will probably also take place this year.

The result of their fee structure is it discourages alternative parties to enter the global IP backbone market when confronted with high prices for transit services compared to leas-

ing a private line or wavelength into the USA, it thus reduces demand for wavelength services as well as rations demand for international backbones crossing US borders.

Operators of backbones that are required to file their capacity use, will shut down most of their capacity in the final days to evade a US$0.93 per 64 kbit/s fee (31 december 2007 rate) that effectively doubles the cost of a 10 Gbit/s. The FCC counts a 2 Mbit/s for 30 circuits and an STM-1 / OC-3 for 1890. A 10 Gbit/s link therefore is counted for 120,960 circuits.

Those who have observed the global 5 minute and hourly Internet Bandwidth Statistics at Internet Exchanges (e.g. http://www.ams-ix.net/technical/stats/) are often aware there is a demand slow down in the Christmas Season.

The supply slow down due to deliberate decommissioning is not easily viewed from usage stats, as international capacity throttling by partially shutting down major links, is not easily measured.

But if you experience a slow Internet the next days, despite the fact that it is the annual slow season and capacity at first sight would be underutilised, you now know the real culprit: the FCC, who's levying of international bearer circuit fees has gone unquestioned as seemingly many politicians think it is a tax burden that mainly hits foreigners and not their constituents, so it is a political godsend tax ...

The result of a change in regulation might be that suddenly a large swath of seemingly unused Purchased Capacity is altered into Used Capacity and we start to learn that it is not the Internet that drives global capacity demand, but International Private Lines and the bizarre habit of the Annual International Capacity Holiday will stop. How big this effect actually is, however will be anybodies guess as statistical data now are deliberately kept under the radar.

We will see if the new FCC wants to tackle the issue instead of further postponing it, like the current FCC who seemed unwilling to grasp the big picture involved with their behaviour.

P.S. It is not surprising that the Verizon, AT&T and Qwest departments (common carriers) heavily protest the proposed alteration of fee structure by VSNL (now Tata Communications) in http://edocket.access.gpo.gov/2008/E8-19899.htm, as the proposed fee alteration by VSNL is tilted the other way out of their favour, obviously to create a negotiation position.

I however doubt whether their regulatory departments have an actual grasp of the amount of capacity demand out there and what is now hold back or shut down due to regulation.

What is surprising to me is that this docket is entirely an intra-industry regulatory discussion, while the obvious burden of excessive fees with an absurd LINEAR accounting is ultimately borne by corporate end users of capacity and firms like Google, Microsoft, Amazon as well as media firms and CDN operators, who built global networks between their major data centers.

Is the bizarre fee structure really that attractive to big market players for keeping smaller guys at bay with potential bypassing and network construction, while it keeps corporate end users sticky in using the more expensive IP-VPN packet services and media companies overusing CDNs and influences network designs?

Vest: Very interesting. Do you have any sense of the historical genealogy of this tariff scheme -- i.e., has it been around in this form for a long time?

Rood: http://www.fcc.gov/fees/regfees.html#whowhat IB (2004 Bearer Circuit Clarification Public Notice):

It opens: PUBLIC NOTICE

Released: July 6, 2004 DA 04-2027

COMPLIANCE WITH REGULATORY FEE REQUIREMENTS BY CABLE LANDING LICENSEES OPERATING ON A NON-COMMON CARRIER BASIS

To sell capacity on U.S. international routes, the operator of a non-common (i.e., private) carrier submarine cable landing in the United States must obtain a cable landing license from the Commission prior to landing and operating such a submarine cable. As a provider of active capacity on U.S.-international routes and as an operator regulated by the Commission, such a non-common carrier submarine cable operator is subject to the statutory regulatory fee requirements administered by the Commission. Since 1995, the Commission has stated in its fee guidance that all cable landing licensees operating on a non-common carrier basis must pay regulatory fees for all international bearer circuits sold on an indefeasible right of use ("IRU") basis or leased to any customer, including themselves or their affiliates, other than an international common carrier authorized by the Commission to provide U.S. international common carrier services.

Vest: Do the statistics you're referring to enable one to definitively correlate apparently oversized purchases with likely users/sellers of IPL services subject to the tariff? I ask because when I was in the business of making (or questioning) such over-purchases, this factor was not at all relevant -- either for the large Internet access provider I worked for or for any other over-purchasing large ISP that I interconnected with. At that time, strategic overbuying was commonplace, but the reasons had more to do with private industrial policies (e.g., peering and interconnection, reciprocal supplier arrangements, etc.) than with public regulatory arbitrage.

Rood: In April 2000, when I published an article "Indicators for bandwidth demand" in Telecommunications Policy, I was questioning the accuracy of the FCC statistics on international circuits with publicly known examples of high bandwidth links into the USA that for some reason did not turn up in the official statistics data.

The main examples I then used where some known high-bandwidth International Private Lines into the USA, which would then materially affect bandwidth statistics. In one of the following issues of *Telecommunications Policy*, Douglas Galbi responded in a note to his article later that year. It is in pre-print version here; http://www.galbithink.org/bandwidth.htm

"[18] As Rood, Hendrick, "Indicators for bandwidth demand," Telecommunications Policy 24 (April 2000) 263-270 points out, U.S. international capacity statistics have significant weaknesses. U.S. international private line revenue statistics are particularly suspect because the revenue associated with international private line connectivity, such as international Internet peering, is not easily accounted for. As Rood notes, the high capacity circuits counts associated with these revenue statistics appear to be inaccurate. Aggregate international capacity figures can be highly distorted by the treatment of intermediate links. For this reason I have focused on aggregate capacity reported by discrete trans-Atlantic and trans-Pacific cables."

I then still thought one of the main sources in the bandwidth and traffic growth fuzz of that day, was the peculiar method of bandwidth accounting by Worldcom/Uunet, as they had counted their growth for years in

bandwidth-distance product but discussed it in the general press as growth in bandwidth and so created a very wrong perception of actual bandwidth usage. I made some remarks about that method of accounting for bandwidth in the *Telecommunications Policy* article.

I then however was unaware of the genuine cause of these non transparent statistics, which concern International Private Lines (the FCC does not track traffic utilisation on packet networks).

I do know now better, that the real source of underreporting was that some newer carriers did not know they not only had to file, but also that the reporting was used to levy regulatory fees on them. Several did discover that as they received a hefty surprise bill from the FCC, others found ways to avoid reporting them.

The FCC's own regulatory fees cause a major distortion not only in international bearer statistics reporting. The impact on actual operations behaviour is to reduce IPL-based (backbone) capacities after Christmas and postpone new activations of backbone bandwidth to avoid the FCC's excessive regulatory fees.

If you are a careful reader, you can glance from those public filings in the docket that this annual regulatory driven capacity shut down is mentioned by tiptoeing, if you want to observe as an outsider the actual shut downs that sometimes slows your Internet connection in the final day of the year, you may need quite some instrumentation or deep digging through route announcements. But that however does not allow you to observe similar practices in corporate networks, it gives you only a view of the public Internet part of international capacity from and to the USA.

As far as I can see back the fee structure is in place since 1995.

Coluccio: I may be all wet about this, but your thesis concerning underreporting aside for a moment, I am left to wonder if the "used" capacity vs. "purchased" capacity figures can actually be regarded as a utilization rate, which we'd normally see represented as a percentage of actual throughput (and overhead) of the total connected capacity. I.e., the 2.1 Tbps used out of a total of 5.7 Tbps "puchased" = a utilization rate of ~ 37%? Or, is this something you've already surmised but I'm missing it?

In a larger context, I'm reminded by your post (which I thought was excellent, btw) of the debate of the 'mark-to-market' rules vs. the actual daily valuations that recently ensued after the financial markets began to unravel.

Newman: Yes I first took "Used Capacity" as that which is occupied by real traffic, as well.

Utilization of 37% is pretty high (given conventional protocols etc.) so I would like to have confirmation of the meaning of Used Capacity as well.

Coluccio: Likewise, I thought 37% was on the high side, but then I considered that the "total" also included a fair amount of ATM and other TDM derivatives, which are more tightly packed.

Rood: The 37% in use by *Telegeography* as a percentage of Purchased Capacity refers here to sold leased line capacity on the activated transmission systems (Lit Capacity) - blocks of lightpaths - on the transatlantic route.

It is all a layer 1 discussion.

Coluccio: Standby capacity used for one-for-one fail-over (contingency capacity use on self-healing SONET/SDH rings, despite some being link-switched) is both purchased, and, during normal operation, "unused". Transatlantic routes have been particularly notorious for this

type of architecture since the late Nineties. How is contingency capacity counted?

Rood: Potential Capacity is when all the submarine cables are fully built out to their limits as foreseen during design.

In practice this means the ultimate design capacity including foreseen major upgrades.

So the Potential Capacity of the TransAtlantic submarine market is the sum of all bandwidths at end of any upgrade cycle of all the systems currently installed in the Atlantic.

A fibre optic submarine cable is typically designed today for a few Terabit/s as potential capacity, that often includes major upgrades, e.g. not only the addition of more wavelengths, but the replacement of optical submarine terminal systems by one more advanced generation of transmission technology at higher bit rates.

The last time many systems "jumped" with upgrades was around 2000 when on several systems the 2.5 Gbit/s wavelengths were replaced with 10 Gbit/s wavelengths. Then if you have replaced all channels with 10 Gbit/s the system has reached its Potential Capacity, as it is typically not possible to design the lengthy and rather carefully engineered submarine systems of many thousands of miles with two generations of improving bit rates.

And a few hours later Rood: There are various types of contract:

1. On the ring shaped systems you buy one leg for the capacity and get the protective capability on the other leg
2. End user leases diverse routes and does protection by itself
3. Arrangements are made with lit but unused capacity that can be activated on an ad hoc notice against fees that can run up to millions per month
4. Restoration capacity is preplanned and preprovisioned, but not activated

See:
http://www.seamewe4.com/doc/Use%20of%20SMW4%20to%20restore%20other%20Cab le%20Systems.pdf

For arrangement 3 all Lit Capacity can be brought to the table
For arrangement 4 there is a payment and a purchase agreement
Arrangement 2 is definitely "In Use"
The classification of the second leg in arrangement 1 is unknown to me. Could be either "In Use" or "Purchased" depending on the priority of the contract. If it is a 1+1 or 1:1 protection, I would classify that as "In Use".

1:N protection on the same cable system is not so much route diversity as it is protection against submarine terminal equipment failure. Due to the multiple failures in cables in the Mediterranean some operators are currently rerouting from Europe to Asia via the Atlantic, through the USA and over the Pacific.

It is an interesting question whether such ad hoc capacity provisioning on a daily basis will be counted this Wednesday for the FCC regulatory fee.

Robert Atkinson: If this is happening as described, once it becomes aware of the subterfuge the FCC is likely to do two things:

1. change the reporting basis so that it isn't a snapshot on a particular day

2. commence enforcement actions against these users who are avoiding their USF contributions through the subterfuge. Because of the growing need for more revenue for the Fund (that's another discussion entirely), the FCC has been imposing fines against companies who are intentionally avoiding their obligations. This sounds like an "opportunity" for the

Commission. Good work for the lawyers.

Rood: No Bob,

That "USF good lawyers stuff" cannot be the case. The USF-fee is based on annual revenues from international leased lines services not on capacities.

There isn't any indication at all that carriers misreport annual revenues on IPL, voice and Internet traffic.

As far as I read it, the FCC's levy is for their cost recovery. So hunting down one subterfuge does not raise FCC-income next year, but redistributes it among the carriers. Regulators tend not to be allowed to receive more than cost recovery. I have seen once such a case, and the overcharged operator won it in court hands down.

Also somehow I get the idea that no one in the FCC does math anymore:

1. The current price for a 10 Gbit/s wavelength/private line London-New York is around US$125,000 per year, about US$1 per 64 kbps circuit 2. The FCC now charges US$0.93 per 64 kbps per year, as a 10 Gbps link contains 120,960 circuits in FCC calculations 3. The combined amount of all 64 kbit/s circuits that the FCC gets reported is 7.5 million. With 6.5 on submarine cables and most today parcelled out in 10 Gbit/s wavelengths.

With all due respect, does the FCC really think that the entire market for international capacity in and out of the USA has declined to a mere US$7.5 million? Did they ever look in what they got reported for USF on revenues in these markets? A short glimpse at construction costs of submarine cable systems and the fact that there are still new system launches?

The reporting of International Bearer Circuits includes all transmission capacity, not only for private lines, but also for voice, data and Internet links.

Miller: I am not sure that the averaging windows would be consistent across the fiber links which may skew the data of purchased/used. 1 min averages being combined with 1 hour averages…. or peak reporting?

I would also be curious as to how the "Potential Capacity" is factored. Is that 10Gbps X64 wavelengths per fiber? Does "Lit" mean at 10Gbps or at any speed?

Based on slide 26, purchased capacity is a 5.7Tbps at the end of 2007. I am not sure of the viability of trans-oceanic 10Gbps or number of wavelengths on _older_ fiber, but my napkin math says 5.7Tbps = 570 x 10Gps links. AT 32 lambdas per fiber, that's only 18 pair of fiber. Lets be conservative and say 16 lambdas per fiber for 36 pair…. I checked my match a couple of times… I feel like I am dropping a decimal somewhere… :-)

Coluccio: If I'm not mistaken, the figures shown in the reference slide are "aggregates", representing the total of each category over all transatlantic routes. Most subsea systems that extend over several hundred km do not possess more than eight working fibers or four fiber pairs due to the challenges associated with powering repeaters far from shore. WDM ratios, however, are higher than you've allowed, with some as high as 64 wavelengths or greater per strand, and some are now boasting the ability to double those, still, through dual-polarization modulation schemes.

Telcos Invade Each Other's Turf?

January 1, 2009 ***COOK Report***: Is this pre infrastructure bill debate astroturf? Looking at what follows, it seems to me that Bennet is breathlessly over reaching in his desire to prove that the monopolists are our "friends"

Richard Bennet: Network World reports that Verizon is over-building into AT&T territory in North Texas, providing consumers with a choice they've never had before. I don't have details yet, but some of your readers certainly do.

http://www.networkworld.com/news/2008/123108-fiber-att-verizon.html?fsrc=netflash-rss

Verizon is setting up a Wild West-style telecom showdown by expanding its FiOS network further into territory traditionally held by rival AT&T, says a new report from Information Gatekeepers. According to IGI, a telecom consulting firm, Verizon's recent FiOS expansion into areas of northern Texas could mark the first time that one carrier has directly competed with another in its own franchised territory for residential wireline Internet services. Traditionally, Verizon and AT&T have competed with each other primarily for wireless voice and data services, as the companies' landline businesses have been dependent on architecture that each company has purchased over the years from the original "Baby Bell" companies formed in the wake of AT&T's breakup in 1984. But with Verizon now offering video, voice and data services over its fiber-optic network in AT&T's home state, IGI says that the telecom industry could be "drastically" changed. In particular, IGI says that Verizon's decision to "overbuild" its facilities into AT&T's franchise areas could spark AT&T to begin overbuilding as well, thus turning competition for building out services into a potential "nationwide phenomenon."

Goldstein: Bennet is not necessarily "friends" with the monopolists, but his gadfly streak shows up in strange ways. Still, this is an interesting story that he passes along. It first popped up a couple of months ago. Essentially, when the Telecom Act was passed, there were seven Bells, and GTE was larger than any of them. The expectation was that they'd compete with one another. Merger conditions for SBC' acquisition of Ameritech required them to compete in some number of out-of-territory markets, or pay a fine. Lo and behold, they did a little token resale competition, and I think paid a token fine, but gave up. BA/VZ and SBC merely competed on who could buy up the most exclusive turf. They had a sort of unwritten pact to never compete on wireline LEC business. Their employees with competitive instincts had the mobile side to play with.

North Texas is a kind of interesting market, though. SBC/ATT has downtown Dallas and Fort Worth. GTE had a band of outer suburbs north of the city, with SBC turf farther north. That area has seen explosive growth, mostly upscale. So the GTE->VZ turf was one of the first to get FiOS. ATT is trying to milk its old copper plant with DSL Lightspeed, which is probably a harder sell. So VZ is setting the non aggression pact aside and slipping across the border into the subdivisions on the SBC side. It doesn't cost any more to pull new FiOS plant on SBC's side of the line, after all; it's still near enough to the head end/CO, and this way they're cannibalizing somebody else's

"access line" count, not their own.

What's really supposed to be odd about the story, then, is not that it's taking place, but that it's taking place on such a small scale, and took this long to happen.

When I was at GTE in 1997 (when they bought BBN), they had a CLEC division, and it seemed logical to me that they could compete head on in a lot more places. GTE after all had a big presence in the Los Angeles, Seattle, Portland (OR), Triangle (NC), and other significant shared markets. But they instead mostly just dallied with resale, then sold out to BA (when I promptly left), and that was that. If VZ does decide to go head-on with ATT, the Los Angeles market looks very promising. But then ATT could also build into VZ turf, attacking the more profitable non-residential markets, so there's risk.

Coluccio: Fred, your trailing comment concerning out-of-area attacks on the more profitable area of commercial customers is, in my opinion, dead on. It's a perfect segue for Verizon, say, to go after the larger enterprise market, and this can be achieved through resale, i.e., through Type 2 offerings at first, i.e., resale of loop facilities, as opposed to Type-1, which would consist of a facilities-based offering.

We're already seeing signs of this through their unregulated arms. VZ, for instance, is already wooing out-of-area government installations and small-to-medium-sized business with its FiOS offering, something they'd be loathe to do at this time in their own markets for fear of cannibalizing their own commercial-grade transport services. Think of the many asymmetries at play here between the respective offerings of VZ and AT&T that would be on the side of VZ, in this case, it if faced off against AT&T's U-Verse offerings in industrial parks and inner-city business districts.

From there, it would be a small task to qualify for offering services to municipalities and high-rise residential (or mixed-use) MDUs and government buildings. The question that still rings loud within me is this:

Does the same unspoken/unwritten non-aggression rules that you alluded to in your earlier reply still apply, now that there are only two?

To expand a bit on what I meant by asymmetries, consider this brief passage from a chapter of Clayton Christensen's book *Seeing What's Next:*

"The more interesting scenarios occur when there are asymmetries-important differences of motivation or skills. Asymmetries of motivation occur when one firm wants to do something that another firm specifically does not want to do. Asymmetries of skills occur when one firm's strength is another firm's weakness."

From:
http://hbswk.hbs.edu/item/4353.html

The Innovator's Battle Plan

COOK Report: In doing FiOS is there a fundamental change in Verizon's network architecture that makes it far easier and more cost effective for Verizon to bring high bandwidth service to the enterprise as well as the home?

Coluccio: I pointed to FiOS only because it is the most recognizable form of optical delivery that VZ is making available today. However, in its current state it is not optimal for enterprise use for most purposes. As a PON offering, it could be made to fit, sorta, but would be better substituted by point-to-point Ethernet.

If the form factor of FiOS's CPE were modified for use by the average information worker (think the size of a

standard wall jack), it would in many ways be suitable, but the analog video component would need to be tossed, of course, since the latter accounts for a large part of CPE overhead, and would be considered overkill (or an unnecessary form of superficial cost) to the needs of most businesses, especially those who've deployed video over IP capabilities.

One aspect of FiOS that is working in its favor right now is that it lends itself to the cloud model, since, by definition, it would be a managed service, hence freeing up enterprise personnel and other resources from the drudgery of repetitive administrative tasks (such as IP subnetting, router configuration tasks, distributed security appliance admin, etc.) to do other things. The downside, of course, is that, in its present state it would be controlled entirely by a single provider (or its contracted partners). And that, in and of itself, detracts from its attractiveness for most large, discerning buyers who have the resources to act more selectively, which is why, I suppose, it's been targeted at the lower end of the markets at first as the home?

The answer to this is situational in many ways. Bulk delivery to neighborhoods differs from how VZ would deliver bulk to enterprise desktops, say. The former lends itself to a dedicated fiber per home, but VZ has already demonstrated that, to deliver a single fiber to every desktop (or to every dwelling within a very-large MxU), it considers doing so overly taxing on infrastructure (and administration). Such scale is more easily satisfied in some of the cases I've seen already with a fiber-to-the-building approach that uses VDSL or ADSL2 in the last several hundred meters.

The economics governing the sizing and siteing of network elements becomes pivotal here. In some instances I can see them going native fiber all the way to the CPE, especially where the future extensibility of throughput may be seen as a competitive issue. In a more diminutive sense, though, such as through the use of xDSL in the final stretches of delivery, they'll take the tried and true approaches of incrementalism, common to how DSL migrations have taken place over the past decade or more, and, I should add, in a manner similar to how AT&T continues to deploy its FTTN to this day.

COOK Report: Is there a well marked dividing line between residential and business services?"

Coluccio: There was a time when I would have speculated that the "well marked dividing line" might be found only in the headers of packets, leaving the job of differentiation to upper layer functions found in the optical line terminals and central office routers and switches. However, after a rather interesting discussion that I had with a friend not too long ago I'm now inclined to think there may in some locales, at least, actually be separate overlays, if not entirely physical, then partially physical and entirely logical. In my friend's case, his adding a business service to his residential account resulted in VZ running a new fiber, complete with a separate CPE enclosure, to the side of his house, while also issuing a separate set of IP addresses.

COOK Report: "If you look at Verizon business services and the old MCI and ATT business services the old ATT," is there not a fair amount of overlap on THAT territory?"

Coluccio: I think you'd need to be more specific. Enterprise-related "business services" aren't as straightforward to characterize in a way that they can be compared to residential architectures. Residential tends to be one fiber per home, and that's it. In the enterprise space there's a greater degree of fragmentation taking place. For instance, an enter-

prise may lease local loops from Provider A on a resale basis that attach to the long haul or Internet facilities of Providers B and C.

As an aside, you may recall that, when the last two mega-mergers took place, the DoJ took measures to ensure that AT&T's (TCG's) fiber assets and VZ's (MCI/MFS's) fiber assets would not preclude competition from taking place on a building-specific or locality-specific basis, if the newly acquired fiber property were in the same building where the acquirer already had a fiber presence. For example, in NY City, where VZ and MCI's MFS were the only two sources of fiber in a building, the newly-acquired fiber assets were either spun out and auctioned to others, such as LEVEL3, or, I suspect in some cases simply redeployed at some mid-span point in the metro to other buildings, or in some cases decommissioned entirely.

Rood: I think the answer [of incursion by one incumbent on to the turf of the other] is relatively easy:

1. Fibre optics (PON or PtP Ethernet) allow for far longer loop lengths
2. They have their head-end serving their former GTE areas for FiOS, so the tech operations is already in place
3. Serving neighbouring areas now served by AT&T/SBC effectively means low additional cost in advertising and other marketing and sales costs.

If SBC/AT&T sticks with VDSL based approach, trying to extend fiber is easy in those complementary areas. It is not in entirely new areas, where there is no such neighbouring area footprint.

It also does not extend to business markets. Counted as percentage of lines/ buildings, businesses only have a few percent of the lines. When you take each SOHO business into account it runs up to about 10% of all lines/buildings. Only in areas with pure business parks and in city business districts a competitive physical medium infrastructure serving businesses only makes sense.

So the bottom line is that this may happen here and there, but you will not observe all out competitive invasion of VZ and AT&T in each other's areas.

Cowen: One thing that is critical to understand in competition for enterprise customers is that competition takes place for multiple sites. It usually takes place against a tender that is issued by an enterprise for its requirements and the enterprise usually has a very good understanding of its requirements.

Typical requirements for enterprise networks supplying voice and (private data) involve router based services (voice and data). These can be supplied by a number of different suppliers: IT players may tender on the basis of integration with the existing IT estate, systems integrated and telcos supply the communications infrastructure and combine WAN with LAN.

The critical point is that competition is for multiple locations. Those with level 1 underlying fibre or other infrastructure use that as part of the offering. Those that do not own such facilities have to lease them from those that do.

This is where a badly regulated monopoly incumbent such as AT+T and Verizon gain huge opportunity. In their ability to offer more coverage over more sites from their own resources they have an economy of scale, scope and network externally advantage. This is well understood and the basis for the need to regulate.

Regulation was supposed to create a level playing field so that those who have to buy-in the multiple site coverage can do so on equivalent terms and conditions and prices.

When we talk about competition the enterprise space is hugely important: the VALUE of the lines to each enterprise is much higher than to a consumer, but most importantly, from the point of view of the public interest, the benefit of communications to the enterprises and their application in the business environment, productivity, efficiency etc is significant.

The lack of effective regulation of enterprise access has been calculated as costing the US 246 thousand jobs.

It was agreed in the European Council at the end of November last year, that continued regulation of access is necessary, and indeed that additional remedies may be imposed by national regulators. Additional remedies will now include functional separation (despite heavy lobbying by German, French and Spanish incumbents).

Coming Changes in the Infrastructure Model?

St Arnaud: "Already researchers are using various wireless techniques such as **quadrature phase shift modulation** to achieve data rates in excess of 560 Gbps on a single wavelength in a DWDM system, and it is expected that data rates in excess of 1000 Gbps per wavelength will be possible soon."

COOK Report: Questions: how soon is soon? James Crowe at Level 3 the other day said something about his fibers being filled. Hmnn - sounds doubtful to me. But while infinite capacity in a strand of fiber does not make telecom free -- there is still op-ex and cap-ex to maintain the network, it seems that it renders the scarcity charging model more and more absurd. As Harvey said the other day: its ROADS.... highways..... The government provides roads and highways as a public good that anyone who operates a vehicle may use as long as they obey universally applicable rules.

Surely the assumption that networks of glass threads are private toll roads must be over-turned and governments must engage in the building of glass roads for the public good? Use will not be free but the system will not be structure on a scarcity model. Open access and common carriage must be determining factors.

Lots of questions including what does it take to get people to recognize this?

Again this is one more reason why I am likely to begin to spend an inordinate amount of time with the material from Harvey on the LHC network and the lessons learned from this global co-laboratory. Questions are how do you get people to understand what is at stake?

Thoughts?

St Arnaud: I suspect these optical technologies will remain in the lab for a very long time as there is simply no market for them. Carriers are still trying to payoff their last investment in 10/40 Gbps systems and in are no mood to make any future upgrades.

COOK Report: Carriers yes..... But doesn't this open a road for governments to build? Its the HIGHWAYS 'stupid' - right? The carrier business model no longer makes any sense??

St. Arnaud: As well the biggest challenge is not backbone bandwidth, but router bandwidth and aggregation. This is where the choke points will continue to manifest themselves for some time yet. And this is all electronics - not optical

COOK Report: Yes indeed - but if government built its own optical interstate highway system, doesn't this augur well for a switched lightpath all optical answer where ever possible? The hybrid networks of Kees Neggers and Cees de Laat? A network for education and e-

science that the carriers will never allow.

Yes? And as Rollie Cole just said consider the case of health care aspects of such a network

So there are really two roles to be discussed -- how broadband can support collaboration and imaging that lead to breakthroughs delivered to the very top of the spectrum, and how broadband can also support delivery of those breakthroughs (and current best practices) to the rest of us.

Cecil: Gordon, et al. good points all, and agreed. I think the question with regard to Government might be slightly more nuanced than simply whether government should purchase fiber optic and get into the information highway business as it were. It may be also - and more immediately since any such mass purchase would require all kinds of political activity - a question of whether in the nearer term we can update the regulatory system. As I said long ago and far away, when the regulators and politicians value abundance, it will be valuable. So far as the present system is constructed, it values scarcity. Rational economic actors within such a system have no choice but to operate within it's bounds. And while particular carriers may merit more vitriol than others, these are matters of degree within a given and expected set of norms which distract from the greater and more important question of creating legal, business and technology cycles so that capabilities are unleashed for profit rather than constrained for profit.

I've been spending quite a bit of time with thinkers in alternative energy, and one person, in particular, who was (or credibly claims to be) responsible for coining the expression "new energy economy" at a time when none of the things that are widely obvious to the greater populace today were all that widely acknowledged. I bring this up because it suggests to me the necessity, acknowledged at times in this forum, of reframing the discussion. In that regard, I would commend some thought around positive vision so that you create a positive feedback cycle where implementation follows principles and principles follow vision. The vision must be compelling and it must be positive.

Along those lines, rather than excoriate networks and network operators who, at high levels, do nothing more than sell goods they are allowed (or tolerated) to sell in the market not unlike any other player, I would offer that all of these providers are valuable inputs into a greater system, but that presently 19th century technological and economic assumptions prevent them from launching and others from offering greater connectivity. While I remain no fan of the 1934 Act and would rather see it extirpated and cast aside, perhaps we recognize present conditions and set forth a vision of the future in a new section of the Act that realigns incentives in ways that encourage any entity - private or government - to build and operate fiber optic systems. In so doing, let's take the explicit and implied subsidies out of minutes and channels, etc. and shift them to glass (I know, I know - this is USF reform and I empathize completely with what that means and what those who are creatures of today's system will want to do with it).

More broadly and more deeply we need a more integrated vision of infrastructure, to which vision I will attempt to make a contribution to in the very near future as I plan to float a paper here proposing an entirely different lens through which we might want to start to view this question.

St Arnaud: I agree with your thinking, but I think a new integrated vision of infrastructure has to be tied into a "zero carbon" economy rather

than a "new energy" economy.

See my various diatribes on this subject at http://www.slideshare.net/bstarn/preparing-for-climate-911-event-presentation http://green-broadband.blogspot.com/ http://free-fiber-to-the-home.blogspot.com/

Coluccio: I take it that you actually mean "conceive in our minds and include in our plans tomorrow's world today", as opposed to create tomorrow's world today. If so, I'd agree, although how one would go about living such a conception today escapes me, except for the pursuit of such, itself.

Cecil: Well, to conceive something is to start. I see it more in terms of a continuum - things conceived in the mind are followed by words and then actions resulting in things; once any thing is built, further thoughts, words and actions are required to keep it running. Accordingly, to agree - as a group - on a starting point - on a vision (and we are very much there in many respects), but to collectively focus that thought energy inevitably leads to words, which lead to actions, all of which are results. In that sense, today's world is continually in flux and tomorrow is always a vision (it never arrives except in the present moment). Thus, at the moment we agree on a vision, all of reality changes whether we can perceive it with our five senses or not. Sounds "out there" but actually deeply practical and confirmed by physics, if not by practical experience.

Accordingly, can we agree that we are in a post-telecom, post-Internet world and the work that needs to be done is recreating not only infrastructure, but how and why we use it. If so, then, that very agreement is the beginning and all that follows, so long as directed toward that end - no matter how prosaic or picayune, will be directed to that greater result. In other words, what activities are done are important but why they are done (or to what end) is primary. If so, then future, past and present unify in each and every activity undertaken as all are directed toward a common vision. Or so it seems to me ...

Harris: [referring to Erik's remarks on the 1934 act above] Maybe it's time to dust off that hoary old notion of "natural monopoly" and make some factual determinations of what parts of the network might still fit that description and what parts don't anymore. Infrastructure that is still a natural monopoly should be operated and regulated under rate of return principles and the rest is unregulated. I know "natural monopoly" is a vestige of the 1934 Act, but I think it's still a very useful concept, especially in mostly rural areas.

Cowen: What is and is not a monopoly depends on the context, and the question being asked.

On context I always think three things are critical; the product characteristics and competitive products depending on demand side drivers, (product market), geographical location and availability of supply (geographic market) and the time factor.

This last point is more than merely whether there is timely, likely, and sufficient response to price and non price factors. Time is both a critical performance characteristic and a productivity factor in many customer applications, but more importantly the response time and transaction costs may mean these factors become irrelevant. To give one example: products and services from different suppliers may look the same, operate within quality and performance characteristics and may be price competitive, but if one supplier (supplier A) can cover all sites with a single offer over existing infrastructure and all other competitors are dependent on supplier A for their ability to respond,

then supplier A can wins by bidding and dealying responses to others that depend on its inputs as critical for their offers. This happens rarely as it is very visible, but less visible and more prevalent is the issue of response in terms of time to repair or performance available to third parties.

The incumbent with the biggest footprint has the ability to discriminate vis a vis others. To keep it simple you could say that Bigfoot is the problem.

This is not a natural monopoly in the sense that the economies of scale scope and externality give rise to unassailable cost benefits.

It is vital that we stop thinking about ourselves as typical customers: customer requirements in certain sectors of the economy, banking, financial services media etc are more time sensitive than others, and the vast majority of consumer internet uses are very different.

Another issue with natural monopoly is that it can be argued that 'SINCE' there will be only one player, everyone else should give up and go home. Personally I think it is critical to set out with the goal of competition where possible and effective regulation where necessary. Regulation should apply wherever competition is not possible; and that is likely to be in very many places and for multi site customers.

For example the recent WIK study on competitive supply of fibre in the EU shows that there will be very little competition in most places in Europe, outside of certain CBDs and, in particular, where multiple sites need to be supplied then the discrimination risk is significant.

Put another way: multi site organisations are at greater risk of monopoly abuse than single site organisations in competitive CBD areas.

Also relevant is the following

BT Report: Economic benefits from providing businesses with competitive electronic communications services

see:
http://www.evua.org/library.html

COOK Report: I have just scanned through the material. EXTREMELY impressive.

Is there any feedback on how these reports have been received on the UK side of the pond?

Cowen: Couple of points to note: these reports were led by us but extensive input and encouragement from major users, and customer trade organisations European Network Users Association etc referred to in the reports.

In terms of impact, they have helped to support the position taken by the European Competitive Telecommunications Network Operators Association, (ECTA)and have been helpful to support the position taken by the European Commission in persuading member states to adopt functional separation as a remedy to the access issue.

(I will forward links to ECTA positions that may be relevant to those on the list separately).

The reports help support the position of major users pretty much everywhere. Interesting equivalent is that in the US the trade associations are all pro incumbent and apart from the Ad Hoc Telecommunications Users Group in Washington, there appears to be no-one covering the same ground with USA data; that is what I wrote to the list about some time ago.

Budde: **The business market is often viewed as a driver of innovations, which are then pushed further downstream once the initial investment costs have been recouped.**

That was the way of things for decades, but the arrival of the Internet

put an end to this business model. It was the 13-year-olds who drove the Internet into the residential market, a fact that took both the telcos and the business market by surprise.

The telcos have never been able to recover. Their whole business model was based on very high margins in the top end of the market (often a few hundred per cent) and these margins slowly came down, depending how quickly the telcos decided the services and products should be provided for customers further downstream

Such a process could easily take a decade or longer.

The telcos were also unable - but, most importantly, unwilling - to adapt to these changes. They believed that their monopolistic position would allow them to continue to force their decade-old models onto the market. This is the key reason there is so much turmoil in telco markets around the world.

Most governments had completely, or largely, retreated from this market during the privatisation processes of the 1980s and 1990s, but the old business models are now beginning to hurt societies and economies and this is forcing the governments back into the arena.

With Internet and broadband access the battlefield is now located around the infrastructure needed for the digital economy. Can the old business model do the job here? Could the enterprise market drive the rollout of national fibre networks?

The answer is no. In general terms the volume size (number of premises connected) of the business market is relatively small - often not more than 10% of all connections. It is therefore unlikely that the business market will be a driver behind the rollout of national fibre (to-the-home) networks.

Nevertheless CBDs and business centres are viewed as priorities in such cases. However, apart from the CBDs of metropolitan cities and a few newer business parks, fibre-to the-premises is certainly not widespread. Without strong government guidance it is most unlikely that the usual top-down approach by the telcos would deliver that infrastructure in any timely fashion.

Where we do see some progress is in situations where enterprises - but more likely - large government organisations (healthcare, education, municipalities) join forces and use their combined market power to set the agenda for new fibre infrastructure.

Cowen: I agree. Schools and hospitals are great examples of enterprises that are adversely affected and likely/ some cases not so sophisticated as business customers to look after themselves.

Cole: The fact that you are asking the question shows how unique BT is in the world of telecom.

I would not presume to suggest technical guidance to potential suppliers, but suspect that the answer has to do with density, value-added for the end-user and the like.

From the economic/social/ political demand side, I think there are two opportunities.

1. Single site enterprises that will get one level of deal if they operate alone, but perhaps a much better deal (that they can then share) if they work together. My example here is Notre Dame, which got terrible bids for Internet connectivity on its own, but once it built a ring that connected it with the city of South Bend and a major hospital, the bids were so much lower than they completely justified the costs, economic and otherwise, of building the rings.

2. Multi-site enterprises, such as school districts, real estate

companies, banks, municipal government, perhaps even fast-food franchises (I row with someone who owns 5 McDonald's outlets), pose an interesting situation. They are the ones, I suspect, that led to the draft that Paul sent to us initially. Clearly someone who has "pipes" throughout the area (whether coax, copper, or fiber) is in a better position to offer a deal than those who do not, unless those "pipes" are limited to pure connectivity and required to be shareable. Thus, in the dial-up days, one could get multiple bids for ISP services.

The situation flips over, however, if the sites themselves are already connected. As Bill St. Arnaud points out, these distributed sites could easily become "nodes" in a system that could cover all the nearby SOHO facilities. He is not the only one to understand this -- Qwest agreed to help the city of San Francisco connected its distributed sites, but the city had to agree NOT to share those connections with non-city entitites, thus cutting off the St. Arnaud possibility.

3. So my bottom line is that yes, if we could get K-12 buildings all connected with high-capacity pipes (almost certain fiber), withOUT a probition on later sharing, we would have taken a big step forward.

In another context, it has been suggested that rural deployment could proceed much sooner and cheaper, if it did not start with fiber to every farm, but fiber as backhaul along major county roads, with high-speed wireless to reach the farms along those roads. In the short run, the speed could be almost as fast, and the fiber could be deployed later.

Alexander Harrowell: We had a very similar proposal to this at one of our Telco 2.0 events. Essentially, the idea is incremental or opportunistic muni-fibre - you put in links between your major traffic sources, taking advantage of whatever RoW and layer zero infrastructure you already own, stick a whacking great L2 Ethernet switch on the end, and after a while you find you have a sizable fibre ring around town, which other parties might find useful.

And you write in the contract terms that third party interconnectors must provide for further interconnection; so the bank that hooked up to the switch in the bus garage has to put one like it in its basement, so the dog wash can as well. In this way you create both the high density access fibre and the MAN infrastructure needed for a further deployment of fibre or VDSL or wireless to the home.

Here's a link: http://www.telco2.net/blog/2007/10/incremental_munifibre.html

COOK Report: Bingo! Alex Outstanding!

I just read the whole "blog" entry and encourage others to do so. The highlight -

"Holistic view of cost

With this model for incremental deployment, you keep rolling until you cover the whole city. It has the advantage that there are no leaps of faith; you simply install links where you have a need for serious bandwidth, or where the cost of telco transit hurts. The EU and incumbents can't complain because you originally built it purely for your own needs - right? The open access tariff is a secondary motivator.

It's only in the later stages of deployment, where the network has to add leaves to reach odd patches of city property in residential areas, that any of the usual problems emerge. But by then, you've already wired up everywhere that's reasonably traffic-dense, and that's a start. It's not restricted to city authorities, either; a group of businesses or other organisations could get started building a shared fibre network in the same way, rather like this very cool L2

Ethernet operator in Wellington, New Zealand.

By the end, you've taken a lot of cost out of public service delivery and private enterprise within the city. So by shifting their tax dollars from opex spend to capex, with a small and temporary investment blip (with quick payback), you end up with cheaper goods and services. The real cost to users isn't the telco bill they see in person; it's the ones they pay by proxy."

http://www.telco2.net/blog/2007/10/incremental_munifibre.html "

COOK Report: Who else besides Malcom Matson?

Harrowell: The article hinges on Roy Gladwell of Connected Real Estate Ltd.

Regarding implementers, Gladwell referred to work being done by some British city councils; I think his slides showed installations in Leicester. Citylink, the NZ shared fibre network I refer to in the post, I think got started this way, but between private-sector partners.

The Internet Archive is doing something similar in San Francisco; it turns out there is some city-owned point to point fibre, and it passes under public housing projects because the city owns them, so the Archive is putting routers in the basements and Cat5e cable into the flats.

Cole: Note that Qwest anticipated this process -- it build conduits to connect City of San Francisco buildings on the express condition that the city was PROHIBITED from sharing the connections with any non-city entity.

It is possible to identify lots of fiber flowing through a given area, especially here in Indiana. It is very hard to get the owner/operator of that fiber to "share it," even when it is used at 0.0001% of capacity or so.

Just a caution -- I like the idea,but the "sharing obligation" is crucial to making it work.

Harrowell: see http://www.connectedrealestate.co.uk/ and http://www.theregister.co.uk/2008/03/28/internet_archive_public_housing/

Cole: The one place where enterprise could play a very effective role is in smaller towns and/or rural areas, but the enterprise and/or its provider MUST be willing to share. The concept is that the big entity does pay for a fiber connection to the wider world, and then shares that connection (perhaps via WiFi or WiMax) with those "near" to it.

A second concept is that the fiber that runs from one network to another (say one small city to another) is "shared" by placing wireless towers at strategic points along the connecting fiber. The small city at the end of the loop is the cost-justification for laying the fiber; the wireless access points along the loop are the way to hit the individual farms and "towns too small" along the way.

So...in general, big entities have not led the way because either they, or their providers, are not willing to "share" the "pipes" going to and from the big entity.

But this is NOT irreversible. Auburn Indiana is again a case in point. The muni had fiber for its own SCADA, but was not sharing it until the alternative was to lose a large employer in town. Once the muni and the community had experience with sharing, and learned its benefits far outweighed its costs, they are now extending that sharing all the way to FTTH.

I work with a wireless provider in town that has set up and manages some 60 plus hotspots in Indianapolis by getting local vendors (coffee shops, restaurants, etc.) to "share" their connection via the wireless system this provider sets up. He has a way to send the signal around the

sponsor's local network, so reasonably secure. The end-user has to register and go through a splash screen, but usage is otherwise free to the end-user. It has been widely popular where approved; but has run into opposition when sponsors are nervous about security or have already cut a deal with a for-pay wireless provider. Unfortunately for we end-users, the convention center and the airport have such for-pay deals. But much of downtown is covered with these free hotspots, all based on sponsors being willing to share, and providers not cutting them off for doing so.

As in kindergarten, sharing is the key.

Coluccio: "So...in general, big entities have not led the way because either they, or their providers, are not willing to "share" the "pipes" going to and from the big entity." Many enterprises would be in violation of federal/state/industry-sector compliance statutes and guidelines governing security if they allowed promiscuous sharing, not to mention falling out of compliance with their own internal auditing guidelines for security as well.

Loosely translated: : Enterprises with the most bandwidth to offer are usually the same ones that are forbidden to do so, hence they are least inclined to give. Where enterprises could be more supportive is in the political realm where, say, a shared private metro ring is being planned, or they could avail their rooftops for community wireless infrastructure without demanding outrageous compensation. Hardly what businesses would do, however, if they can garner extortion rates from cellular carriers for the same roof-top or building-side footprints.

Cole: I agree with all this, although it is my understanding that technical solutions to the security questions are possible. My local wireless vendor claims that it can in effect send its signals "around" the sponsor's network, not through it; and one could imagine joining the fiber physically just outside the big entity's perimeter.

Rooftops and outside walls have all sorts of regulations now -- one could imagine some "required sharing" just as we are now imagining "required greening" of such spaces with plants, etc.

So, although the problems are real, they are not necessarily unsolvable, with carrots or sticks or both.

Coluccio: Understood. It wasn't my intention to rain on your ideas (although, admittedly, I seem to be taking a contrarian view here lately :-)

Interestingly enough, my experiences in dealing with community reactions to clients' private microwave systems (and more recently public carriers' cell towers near schools and playgrounds), suggest that, even beyond the issues I mentioned earlier, many residents present impediments to such initiatives on aesthetic and fear (of radiation) grounds alone, a la NIMBY satellite dishes and cell towers, respectively. All of which suggests to me that wireless "broadband", in the many forms we've been discussing here ranging from grassroots to top-down, could certainly use the help of local government.

Cowen: I think the concept of 'sharing' depends on the question of who owns the asset.

When a big customer wants fibre to the business park, and an infrastructure supplier lays the cable and does the installation, it will typically identify the potential for incremental additional sales for use of that link at that location and the geo sales along the route.. So the deployment to the business may well help roll out for business and other customers.

The business case may vary from a field of dreams model (we will build it, they will come) that was popular in

the late 1990's CLECs, through to more detailed market Analysys by user by building. The point is that the infrastructure is shared across many customers in most cases of commercial provision. The municipal build model is a variation on this except that the builder or the prime contractor is a non commercial municipality pursuing a public service remit.

Where the local municipality is using public funds then, there are often serious questions to be resolved about whether the state is intervening and pursuing commercial goals that should be the province of commercial operators. This is not completely prohibited but in Europe we have rules that restrict state intervention to situations of market failure, (known also as solving the Heineken problem; reaching the parts that others cannot reach).

I don't think there is any lack of willingness to share; I think the question is what are the commercial terms on which the asset is or could be shared. Maximisation of return dominates the thinking of badly regulated local access players. You often find that CLECs then sit under the shade of the price umbrella created by the prices from badly regulated local incumbents.

Municipalities may have different business plans, different credit positions to cover different costs and different objectives so their propositions may be very different from the typical CLEC or telco incumbent.

One thing that might be promoted is longer term arrangements between municipalities and customers where the credit status of the state is of benefit to customers and commercial operators alike (in many cases the municipality will not build itself but subcontract to an operator to build and operate on its behalf). Longer term deals that fairly apportion risk and provide revenue assurance to builders may even bring costs down as well as help finance the build.

Map the Fiber

COOK Report: Tim, interesting and worthy practice you describe. Is this common to the continent as well? How often is it done in the UK?

Frank or anyone else - is it done here in the US? Not often I'd bet.

How feasible might it be for someone to offer a web site that could encourage this sort of thing? In other words using GIS and available maps and known fiber, if I am company x or business park developer y, why couldn't I go to a web site that would show me the backhaul issues and interconnect possibilities from my location as well as possible business along route that might want to partner?

Now is all this stuff a trade secret for everyone involved? Proprietary? Do not disclose?

Rather than give a bazjillion bucks to the incumbents :-(... why could not the new administration endeavor to set up this kind of web site and think of some economic sweeteners to encourage use?.com

Cole: Our experience in Indiana is that the exact location of such "potentially shareable" fiber is treated as highly proprietary and a whole variety of reasons listed, started with the terrorist/vandal threat, pressures for sharing, competition, et al.

Also, as Ken Miller and others have pointed out, there are technical, legal, and philosophical barriers to sharing, especially if the risk to the sharer goes up and the benefits are abstract, such as "helping the community."

Also certain providers would rather sell (and get paid for) running from the home office out each time, whether they merely extend from an existing customer site or not.
Homeowners are "forced" to share, in a sense, due to the

PON architecture; but like the choice between a party line and a private line (how many of us are old enough to remember party lines?), the tendency is for all parties to choose a private line if the buyer can buy, and the seller can get a higher price than the party line.

Clark: I'd encourage list members interested to have another quick glance at http://broadbandmap.govt.nz/map/ I think with only one exception we have street level or corridor level supply information for all NZ suppliers [Telsta clear being the exception - they provide merely area coverage denoted by polygons]. [Oh, and we did have Vodafone's 3g coverage map also, but that seems to have gone.....] have a play - zoom into a geographic area and select up to 3 suppliers. Telecom NZ has the largest coverage [denoted as Chorus Fibre] - regional players only in their town. The operational separation of our incumbent was key to getting this progress, but also subsidised investment in regional and urban fibre networks to which conditionality of open access at duct or fibre level was mandated.__I think the economic constraints will accelerate traditional thinking to move towards Telco2.0 business models.__As for terrorist threats - it would be very easy to disrupt a telcom infrastructure even w/o knowing exact locations of buried cables [which, lets face it, the owners often lose or don't know precisely]

COOK Report: Rollie, I'd guess you are precisely correct. Yet if I remember how fiber got started in Holland (according to Kees) it was that the Dutch government passed liberalization laws circa 1998 that encouraged sharing. In our country now things are going to HAVE to change. I am simply wondering whether it is impossible now to do something tat encourages a more open less proprietary attitude?

It seems to me that either the feds better build or, if the private sector doesn't want that to happen, it better be open to taking other paths than totally proprietary beggar-thy-neighbor stealth. Can anything be done to increase the level of understanding of why this matters?

If Holland could do it, under Obama why can't we?

To Donald Clark - Bravo Donald. How did this happen? I believe actually that i looked at this once before.... it must have left something in my subconscious that caused the idea to surface. The NZ gov't mandated this..... right? Was there much push back? How much is it being used?

Kenneth Miller: Along these lines, I had been raising the question with the Indiana State GIS system. (Take a look at: http://www.in.gov/igic/ click on map tab). I ask why they do not track Fiber, Central Offices, LEC copper plants (or even boundaries). Part of the point to some of the demographic data (education levels, etc) is to help build business cases for companies to move into the State for economic development purposes. Wouldn't a clear picture of Fiber and Broadband concentrations (and vacuums) be a good way to convince high tech companies to locate to a bandwidth rich State?

The terrorist excuse is somewhat limited given that the GIS system can black this out for non-registered users. This is how they handle high voltage transmission lines.

I think this is a little ironic given that because City, County, and State government is granting right-of-way and in most cases is recording this in the individual municipal GIS systems for civil planning purposes. Although this is limited to conduit and does not indicate numbers of fibers or functions.

Speaking as a corporate customer that has purchased Carrier fiber, dark, and build my own private routes, getting maps out of carriers is next to impossible. But maybe I just have never

spent enough money with them.

Harrowell: The UK created a national database of mobile base stations, which was available on the Web. Since 2005, however, the carriers fell out with OFCOM and stopped updating. Apparently T-Mobile never provided any data anyway. And (as I learned to my annoyance last week, working on an LBS hacking project) there is no API of any kind and apparently no staff who check the webmaster inbox.

Perhaps it had something to do with the project being a sop to the "electro sensitivity" lobby and their science dodger pals on the national press.

There is apparently a study going on into duct/layer-0 sharing at OFCOM, but we haven't seen the results yet, and I suspect it will be more along the lines of "it would be nice" than "you've got to do it" or "you've got to do it and you send the SOAP document to ductsharing.ofcom.org.uk :8765".

Clark: No – the New Zealand government didn't mandate it, just used its influence to encourage. I think a growing recognition that no one can afford the capital required for the country's FTTx upgrade alone has moved people. There is still some reticence from the remaining vertically integrated providers [Voda / Telstra Clear]. The work started with some mapping we were doing to understand how to help people connect to KAREN. The State Services Commission were working on some related activity and picked up the work and drove it to a web-based service [two external consultants were key]. We [the SSC] have some APIs in beta that will allow read and write / update access to the underlying data tables without going through the web interface. Demand side information is still lagging - some consumers and business have started to add and map their demand points but hopefully the APIs will increase this.

COOK Report: Politically this might not be really easy but WHY in heck can't someone have the guts to say that since the stimulus money coming up is supposed to have something to do with strengthening the national economy via INFRASTRUCTURE improvements, you don't get anymore unless you agree to use it in this way. IE by being open about where you nets are and by being open to interconnection! The days of Bush giveaways are GONE. Dream on?

If the taxpayer is stuck with the bill, it is time for Barack to demand that the tax payer get something in return!!

Estrada: Regarding the mapping...

When we did the mapping (GIS) of the north coast of California, it became very easy to see the 'digital divide.' Even though most of the locals thought they knew stuff, the fundamental issues of lack of backhaul (from the big carriers), lack of competition, geography/population/business density issues and distance to an Internet/fiber backbone became very obvious when mapped.

Tons of maps are here for those inclined to peruse them: http://redwoodcoastconnect.humboldt.edu/ There is an interactive GIS version and a bunch of static PDF versions.

WRT the stimulus package, I truly hope that it can be used to drive backhaul builds that allow the smaller WISPs or the cell carriers to provide decent broadband solutions to the areas where there is no ROI for the large traditional cable/telco carriers. Or it can be used to provide ICT skills to the unskilled that live in well-served broadband areas. Or provide a stimulus to build the transit exchanges/community exchanges needed to provide useful competition for under served areas.

Book Reviews

Photoshop Lightroom 2 Adventure

January 14th, 2009 by Gordon Cook

The Photoshop Lightroom Adventure: Mastering Adobe's Next-generation Tool for Digital Photographers is my absolute favorite of all the O'Reilly photo guides - although the Rocky Nook Press books are themselves outstanding. Probably a year and a half ago I requested a copy of the first Lightroom Adventure book. It was so extraordinarily well done and such a visual treat that I shelved my Aperture software and bought a copy of light room. I love it and find it utterly indispensable. However, without a manual, it is pretty much unusable. The O'Reilly book folk sent the Adobe photographers to Iceland and produced a volume of both great beauty and utilitarian value.

Last summer I paid for Lightroom upgrade and as soon as I saw Photoshop Lightroom 2 Adventure, I pleaded for my review copy.
It is a gorgeous and absolutely indispensable book for all Lightroom users. This time the Adobe crew was sent to Tasmania. It is a source of pleasurable frustration as I am trying to do my newsletter work and digitize some 15,000 color negatives from Russia and the Himalaya, that the time spent inside the book has not been nearly as long as I would like.

I have used it to decipher some of the post-processing tools – and there are more of them in this version. One of the ones that I look forward to is the graduated filter described on page 110. I remember this from the 1990s when the physical filter was called a graduated neutral density. I bought one for my old film camera and tried to use it but with uncertain results. The example in the Lightroom text is one of a Tasmanian beach and sky where the contrast of light in the sky and shade on the beach yields a situation where either the sky is washed out or the beach too dark. The filter however produced such stunning results that the photograph is used on the cover of the book,

A few two-page spreads are used to advantage to show what filters can do. The left-hand page without the filter – the right hand with. Negative clarity on page 220 and 221 and the very subtle Tone Curve Adjustment on 228 and 229.

The book is a visual and artistic delight from cover to cover. Mandatory for the Lightroom user and so well done as to make anyone who thumbs through it want to have and use the software.

The Canon EOS Digital Rebel Companion

January 14th, 2009 by Gordon Cook

Being a considerable fan of O'Reilly photography books, when I saw the blurb for the new Canon EOS Digital Rebel Companion by Ben Long I requested a review copy. Having purchased a SIGMA SD 14 almost 2 years ago, and being less than happy with its operation, I decided to get out another somewhat less expensive digital SLR as a spare camera for a two week trip to Greece last October.

With the O'Reilly book in hand I bought myself the Canon EOS Digital rebel. What I would really like is an O'Reilly book for the Sigma which I now understand I must update to the latest level of firmware. Once I do I am told that I should expect much better behavior - earlier problems have been short battery life and after the first three or four pictures the shutter refuses to fire.

But on to the book which for

the most part is quite good. One of my problems is that I am still self-employed for things other than photography. Therefore I tried to steal some time on the way to the trip. At the Newark airport in October waiting for my flight to Athens. I poured over the book and the new camera. I quickly had a considerable disappointment. The book was unable to tell me how to do the first and major thing that I wanted to do with the new camera.

Something short and simple. Namely set the camera to take pictures in RAW format. The information simply wasn't there. It is discussed on page 138. The camera menu combines RAW and JPEG with three different file sizes. The text on page 138 guides you to the quality menu and then informs you that "from this menu you can choose from three different image sizes: L, M and S. Second for each size you can choose from two different levels of compression. You can also choose to shoot RAW. Informative, but only in the most general way. I tried and tried without success to get the camera adjusted to RAW and the largest size. In frustration I got out the manual that Canon provides and, using it, found out how to make those settings quite promptly.

The purpose of the book seems to be to guide the camera owner whom the author assumes is probably new to photography and certainly to digital photography in the ways to use the camera and think about the aesthetics of what he or she is doing as well. Nothing wrong with that but for my purposes not the ideal.

The book is well-designed and well written and attractively laid out. It covers more advanced areas such as white balance and gray cards. As well as the understanding and use of the camera's histogram. I suppose however what I would be most happy with would be a book that would be both a definitive guide to the workings of the camera as well as a tutorial about how to use it and think about why one uses it.

I am pleased with the camera itself. It survived the fall from shoulder level with no ill effect. The battery seems to last for ever without needing to be recharged. The image stabilization gave me two or three good pictures in low light conditions inside a monastery. The 18 to 55 mm kit lens that comes with the camera was acceptable but I really missed the telephoto capability of the 18-200 mm zoom on the Sigma.

Executive Summary

Needed: A New Communications System

We examine the politics of the broadband portion of the stimulus bill.. Mark Cooper proposes a smallish targeted effort aimed at giving **people access to the Internet on open local networks managed by cities and counties**.

"We envision a community-wide fiber network linking all local government buildings, schools, and libraries. The service would be anchored by local government. Non-mobile communications flow over the fiber network. Mobile communications flow over the fiber network to a WIFI/WIMAX wireless network."

In addition to the physical network, Mark also builds in a role for enhancing community social infrastructure and education. "The stimulus package can be used to create a team -- an "E-Corp" -- to train community members in digital communications and digital skills. They can retrain unemployed workers with digital skills to become local tech support. These activities foster the skills for a more competitive work force."

Cooper goes to the heart of the issue when he writes: "The raging debate over how to define broadband for purposes of "special" treatment in the tax code is a dead give away that stimulus spending directed at the big communications corporations can easily turn into corporate welfare. The corporations will use the tax breaks to pay dividends, increase executive pay, or fatten the balance sheet. The way to avoid this trap is to direct funds to local governments and community-based organizations.

This is also the ideal moment to redefine what government can and should do for the people. Providing for the basic means of communications -- paving the streets and building the on-ramps for the information superhighway -- are proper local government functions. The big communications corporations can be hired to dig the trenches the way contractors bid on road and bridge projects, but the people should own the networks and should build the basic communications network that all households need. The private sector can still build its gold plated, hundred megabit network, but it will do so only if people are willing to pay for it. City streets and county roads are open to the least expensive compact car and the most expensive Rolls Royce providing access to basic services for all."

While mark's essay was greeted very favorably discussion centered on whether the last mile should be wireless rather than fiber and how to keep the incumbents from reasserting control.

When asked about the use of wi-fi and wimax, Mark responded: **This is a policy process that is unfolding at a ridiculously rapid pace and I must speak a language that the policymakers understand. My goal is to crate an environment in which we can allow the best technologies to meet basic needs to prevail. You've identified lots of possibilities, none of which are contemplated by the incumbent communications companies in a meaningful way. We use the word WIFI/WIMAX as names that policymakers and the public may actually recognize, but they stand for a wireless solution to the**

basic service for the digital age.

Once the local/middle mile backbone is built and out of the control of the incumbents, the superior last mile technologies will have a much better chance of prevailing. If we can use a community-based approach to basic service and get to 10X10 mbits for mobile and those who don't have a need for or cannot afford a the 100 mbits service, we will have liberated the communications network from the tyranny of the incumbents.

On the following day (January 15th) Mark wrote: The first draft of the broadband bill is out and it contains only $6.5 billion for rural areas. This is a massive victory, since we have stopped them from spending billions to further entrench the incumbents, especially if we can liberate some of the money for a community-based wireless approach, which is vastly superior in rural areas.

Paul Budde points out that if we can build some advanced networks at the edge it is worth doing because once the incumbents finally realise that they will have to come to the party they - in most situations - will use their own last mile connections at prices that make it impossible for wireless operators to compete. The natural infrastructure monopolies will always make it nearly impossible to compete on infrastructure. But one could argue that if we have forced the incumbents to play the open network/competition game and we do end up with low prices, open networks and competition between services (not infrastructure) that we have reached our goals. In the end infrastructure is a utility and I don't have a real problem to leave that in the hands of the incumbents (once again based on open networks, etc).

Harold Feld closes with an eloquent call to action: "There is a tide in the affairs of men, some shakey guy said. And also a good deal of inertia, I will add. Like the Hobbits of the Shire, we lived so well so long most of us forgot that getting justice and changing attitudes is a long fight measured in *years* -- against a well financed opposition with infinite patience and operating on multiple levels. Small wonder that, as years passed and people were prosperous, that the tide flowed with the incumbents and the few reformers and opponents found it hard going indeed.

But that tide is turning and the Shire is rousing. There is an interest and an energy directed at the management of policy not seen in far too long. It is for the most part still unformed -- more an allergic reaction to the last 30 years of free market triumphalism than an organized movement. But it is a real current looking for direction, and it lies with us whether we shall ride it and shape it or whether we shall allow others to but rocks and shoals in its path."

Symposium Discussion: November 18 - January 18

Who Should Provide Infrastructure? p. 15

Herman Wagter: in the Netherlands the physical network is being separated from production, sales, consumption and so on. The physical transmission network is a public utility, the rest is left to the market. The same view is actively supported by the EC. We do not privatize our road network because there is a massive level of innovation in our cars, what we can do with it and so on.

Fibers are conduits for light. The best fiber does....nothing (but guide the light). Only when you start pushing light through the fibers, modu-

late the light, and convert it into information it becomes a telecommunications network. You can innovate as much as you want in pushing and modulating light, it does not affect the fiber at all.

The fiber is a utility.

Cecil: **Functionality is neither minutes nor capacity. It is technology agnostic. Rather it is simply what people do with available resources. And, as you point out, more rational use and allocation of resources makes for better functionality. Lots of functionality rocks; fighting over minutes, pipes, subsidies, universal service for this silo or that silo, by contrast, sucks enormous amounts of air.**

Fiber Unbundling in Holland p. 19

van der Berg: OPTA will regulate the offer done by Reggefiber and KPN as it would with a copper network. So there are protections against price squeeze, discriminatory pricing and excessive prices. Nothing on delaying tactics, which is a bit of a pity as I would think that KPN's All-IP plans have already sown Fear Uncertainty and Doubt in the investors community and this plan just ups the ante.

From the official announcement: KPN will continue its investments in FttH as announced earlier in 2008. However, KPN does not yet have the intention to roll out FttH on a large scale in the Netherlands. KPN and Reggefiber will assess all plans based on timing and location, considering whether new investments can be justified commercially and taking into account the regulatory framework. KPN will proceed with its current FttH pilots in five Dutch cities.

Vincent Dekker's analysis: So unless KPN is putting out a statement that they will definitely not proceed with or decide to any roll out anywhere else before June 2009 (which they won't put out) I don't think they are putting the brakes on anything. They just want to make it look like that.

In the meanwhile, their joint venture with Reggefiber (Glashart) is signing up new fiber customers by the thousands in more then 50 cities and villages all across the country.See my map of The Netherlands here: http://www.trouw.nl/nieuws/economie/article1905291.ece/Glasvezelnet_werft__snel_klanten_.html

BIT Torrent Switches to UDP p. 22

Bittorrent declares war on VoIP, gamers http://www.theregister.co.uk/2008/12/01/richard_bennett_utorrent_udp/

Fred Goldstein: what Torrent is now doing is bypassing the flow dynamics of the Internet and simply blasting away, like a DDoS attack, towards the destination. It won't slow down when TCP applications do. So loss rates at the NAPs in particular will rise. It's thus a social disease, not one confined to its users.

Harold Feld; The critical thing from a public policy perspective is that certain responses in an unregulated market are highly probable. If one subscribes to the theory that the result from the unregulated market (to the extent such a beastie actually exists) is inherently the best result because Coase promised that the market will always reach the most efficient result, then fine. But if you actually want certain outcomes, then one needs to accept that "the market" will not provide them without regulation. snip

If you actually want a result, like a functional internet that conforms to today's existing expectations on openness, then you need to ask yourself how that will happen. You also need to ask what your error correction will be when it turns out you made the wrong choice,

because I stone cold guarantee that any system involving human beings using imperfect information to predict a dynamic future will include wrong choices. And for all you Libertarians out there, failure to take action for fear of making the wrong choice in an imperfect world is as much a choice in this context as the decision to impose a rule or enforce one. This whole "do no harm stuff" is simply code for "take no responsibility and pray it all works out" (with Plan B usually being "no matter what actually happened or what we previously predicted earlier, convince everyone it really _did_ all work out").

Effective public policy must be sufficiently robust to survive irresponsible actors. Irresponsible actors are an utterly predictable set of actors in a field of stakeholders this large. Blaming the actors for being irresponsible is as much a waste of time as blaming incumbents with market power for exercising market power absent regulation preventing it.

Ripe Approves Private Market for IPv4 numbers pp. 30

Note to readers: please start with the summary article on pages 15-18 above.

RIPE adopts reallocation on December 16 2008. Consensus has been reached, and the proposal described in 2007-08, "Enabling Methods for Reallocation of IPv4 Resources" has been accepted by the RIPE community.

The related RIPE policy document is now updated, published and can be found at:

http://www.ripe.net/ripe/docs/ripe-441.html or http://www.ripe.net/ripe/docs/ipv4-policies.html

Further implementation details of this policy will follow soon. The proposal is now archived and can be found at:

http://www.ripe.net/ripe/policies/proposals/2007-08.html

COOK Report - added on January 25 2009 during process of editing for publication: *In one sense what RIPE has done is pushed the allocation decision out to a local Internet registry in other words the ISP. The ISP uses its assigned numbers to connect its customers. If an ISP has a bloc of numbers it is not using, it can now transfer them to another ISP within the RIPE region (presumably for money) and that ISP can use the bloc to connect new customers. The policy pushes the allocation process one level further done and leaves the details of the execution to the decisions of the members giving reallocation power to potentially 11,000 existing local registries that is to say - local ISPs. To participate in the reallocation market one must already have an assignment and an ASN number.*

The RIPE policy goes on to add "LIRs that receive a reallocation from another LIR cannot re-allocate complete or partial blocks of the same address space to another LIR within 24 months of receiving the re-allocation.

The RIPE NCC will record the change of allocation after the transfer. Please note that the LIR always remains responsible for the entire allocation it receives from the RIPE NCC until the transfer of address space to another LIR is completed or the address space is returned. The LIR must ensure that all policies are applied.

Re-allocated blocks will be signed to establish the current allocation owner.

Re-allocated blocks are no different from the allocations made directly by the RIPE NCC and so they must be used by the receiving LIR according to the policies described in this document."

The policy creates an entirely new level of internet authorities that must administer their routing allocations with the same care as the much larger authority.

Vest: **The fundamental break that IPv4 exhaustion/transfers represents will not affect routability or even reliable routability per se -- not immediately anyway. Rather, it will affect who is capable of being/becoming a credible provider of (mostly) reliable routing services for large numbers of customers. If you get IPv4 the old way, you may be able to join that club. If you are not lucky enough to be a member of that group, then you can still become a reliable customer of more-or-less reliable services provided by one of the members of that club.**

Editor – Tom points out the attractiveness of the current situation of the exhaustion of IPv4 numbers for "facilities owners" at the national level, when later he says: for those facilities owners who might be interested in (re)establishing a more modest territorial monopoly, the disappearance of an effective bypass technology [routable IPv4 addresses] might not look like a such a bad thing at all. Better still is no particular effort is required to make it happen -- just do nothing different!

Editor – the address exhaustion problem creates a creates two unequal classes of ISP whose economic playing fields are fundamentally different.

Vest: The "increasing cost and uncertainty in obtaining IPv4 addresses" will never -- can never -- materially affect those that inherit IPv4 today from the RIR allocation era in the same way that it will affect everyone that comes after. The former will always enjoy choices about what kind(s) of addresses that they want to use to grow -- use NAT and private IPv4 (or NAT and IPv6) and end users everywhere pay a complexity premium, or use public IPv4 and the provider (and every other provider) pays a IPv4 transfer premium and a routing system cost premium. Non-incumbents won't be choosing between these alternatives; they'll be paying for both.

What I am saying is that once IPv4 becomes a unique, non-substitutable commodity -- the market for which is completely locked up by current incumbents -- the balance of incentives that drive technology development will tip toward advances that preserve IPv4's artificial value as a bottleneck input, and away from advances that would tend to eliminate that value. Multiplexing advances like carrier-grade NAT are a perfect example -- they make ownership of IPv4 even more valuable. Once providers start bankrolling and deploying products like that in earnest, there's no way that they're going to want to reverse course and support steps that eliminate that bottleneck. And so we'll be stuck with a new non-bypassable protocol-level bottleneck forever --- or at least until TCP/IP is supplanted.

Cole Summarizes

Cole: Let me offer a few predictions, mainly to see if I understand.

1. Nothing much happens, as Tom Vest suggests, until we are "out" of free IPv4 addresses (the historical pattern of waiting for rain rather than fixing the roof in advance).

2. Once out, three things start to happen, at varying rates of speed and with various degrees of success:

2.1 a "market" develops in

IPv4 addresses, with them that has gets -- the more you have, the more you can both buy and sell

2.2 technology to "extend" the use of those you already have (multi-level NAT et al) now gets really serious, even in areas where it was not favored before.

2.3 technology to "get around" the limit (IPv6, something Tom might like <grin>, et al.) begin to receive serious attention, but may or may not begin to be used

3. the whole system is much more chaotic for a while, perhaps a very long while, until and unless a less chaotic system begins to gain traction

4. **Although collectively we might have the technology "smarts" to reduce the time spent in steps 2 and 3, we lack the institutional/social "wisdom" to employ it. (This appears to me to be the crux of Tom's point, but I am happy to be corrected.)**

Vest - **To Rollie Cole:: Very well summarized.**

The only thing missing are the interconnections -- i.e., 2.1 leads to 2.2 lock-in, and as 2.2 continues, 2.3 becomes increasingly impossible -- at least any version based on any form of TCP/IP that is currently available or under development (AFAIK). So what might look like a natural, evolutionary process is really a giant leap of faith -- not into the unknown, but rather into a known bear trap, but with faith that it won't bite, this time, maybe...

Finally **Harold Feld** (a List member) on his blog wrote a provocative summary: http://www.wetmachine.com/item/1428

RIPE Makes Me Vaguely Uneasy By Creating Legal Market For IP Addresses.

I suppose my real problem is **that I just haven't dug into this area enough to really have an opinion. But then again, so few people have - which is part of what makes me uneasy. Few things rival IP address allocation in both importance and breathtaking, mind-numbing technical snooze-inducing power.** This makes it either the ideal laboratory of exciting new ideas or a veritable Devil's playground of possible mischief.

I do know that the longer the transition from IPv4 to IPv6 goes on, the harder it will be to do. Creating a class of actors invested in IPv4 as a valuable asset rather than simply as a question of cost avoidance will not help.

Does what happens in Europe stay in Europe: I'm glad only one RIR has decided to take the plunge. But how long will the other RIRs wait before opting for this approach? We had credit default options and subprime loans for many years before the collapse. Countries and institutions that initially rejected these instruments and investment vehicles as too risky succumbed to the constant pressure from investors eager to "get in on the action." Until last year, countries with conservative investment laws looked like fuddy-duddy dinosaurs preventing their citizens from getting in on the good times. This year, not so much.

I am worried that we will see a huge rush to declare RIPE's policy shift a "success" before we even know what success would mean here. I expect that "success" will be measured in something really stupid and irrelevant, like total number of transactions as compared to assignments. I also anticipate that as parties begin to make money from the sale of address blocks, similar parties in other regions will begin to lobby for these changes.

Without having some framework to figure out what the Hell would constitute "success" or "failure" of RIPE's market experiment, we will almost certainly see it become the dominant model. Why? Because it is the only model, and its supporters will loudly proclaim its success. With no metrics to define success, this will look very convincing.

Editor's Concluding summary has become an article found on pages 15 to 18 above.

International Bandwidth Pricing and the FCC

Hendrick Rood: When I went searching for some international bandwidth figures I bounced into this nice presentation by *Telegeography*

http://www.ptc.org/ptc08/participants/speakers/papers/SchoonoverFinalSlides.pdf

Alongside this Executive Summary it provides a nice overview

http://www.telegeography.com/products/gb/pdf/Executive_Summary.pdf

But it also gives you one pause:

Why is actual Purchased Capacity on lit cable systems (see slide 23 for Transatlantic) more than twice the amount of Used Capacity? It does not make economic sense in a market rife with overcapacity and with falling lease prices for 10G wavelengths to purchase capacity two years in advance and leave it unused.

The answer is, Purchased Capacity on lit cable systems isn't more than twice Used Capacity, but Used Capacity, in particular the International Private Line part is carefully underreported.

The cause of this is the FCC, who thinks it is smart to levy a regulatory cost recovery fee based on a LINEAR count of 64 kbit/s circuit equivalents in use.

Telcos Invading Each Other's Turf

Goldstein: North Texas is a kind of interesting market, though. SBC/ATT has downtown Dallas and Fort Worth. GTE had a band of outer suburbs north of the city, with SBC turf farther north. That area has seen explosive growth, mostly upscale. So the GTE->VZ turf was one of the first to get FiOS. ATT is trying to milk its old copper plant with DSL Lightspeed, which is probably a harder sell. So VZ is setting the non aggression pact aside and slipping across the border into the subdivisions on the SBC side. It doesn't cost any more to pull new FiOS plant on SBC's side of the line, after all; it's still near enough to the head end/CO, and this way they're cannibalizing somebody else's "access line" count, not their own.
What's really supposed to be odd about the story, then, is not that it's taking place, but that it's taking place on such a small scale, and took this long to happen. [snip]

COOK Report: Is there a well marked dividing line between residential and business services?

Coluccio: There was a time when I would have speculated that the "well marked dividing line" might be found only in the headers of packets, leaving the job of differentiation to upper layer functions found in the optical line terminals and central office routers and switches. However, after a rather interesting discussion that I had with a friend not too long ago I'm now inclined to think there may in some locales, at least, actually be separate overlays, if not entirely physical, then partially physical and entirely logical. snip

[Editor: the answer seems to be that there are so many

shades of gradation as to say the answer is "no."]

Cowen: One thing that is critical to understand in competition for enterprise customers is that competition takes place for multiple sites. It usually takes place against a tender that is issued by an enterprise for its requirements and the enterprise usually has a very good understanding of its requirements.

Typical requirements for enterprise networks supplying voice and (private data) involve router based services (voice and data). These can be supplied by a number of different suppliers: IT players may tender on the basis of integration with the existing IT estate, systems integrated and telcos supply the communications infrastructure and combine WAN with LAN.

The critical point is that competition is for multiple locations. Those with level 1 underlying fibre or other infrastructure use that as part of the offering. Those that do not own such facilities have to lease them from those that do.

This is where a badly regulated monopoly incumbent such as AT+T and Verizon gain huge opportunity. In their ability to offer more coverage over more sites from their own resources they have an economy of scale, scope and network externality advantage. This is well understood and the basis for the need to regulate.

Map the Fiber

Pages 77 through 80 above cover efforts to map the existence of fiber as a guide to the feasibility of doing cooperative builds as in Canada a decade ago or at least determining whether and how other customer might be brought into existing plans and make those plans more viable,

Contents

Building a New Communications System for America at the Grassroots Level

A Vibrant, 21st Century Community Communications Ecology Requires Physical Infrastructure, Human Skills and Social Tools

IPv4 Numbers to Become Transferable and Consequently Property

A Introductory and Interpretive Essay for the Symposium Discussion on Pages 30-63 of This Issue

Symposium Discussion November 18 2008 - January 18 2009

Who Should Provision Telecommunications Infrastructure?

Government or Private Entities?

Fiber Unbundling in the Netherlands

KPN/Reggefiber Offer Fiber for Twelve Euros per Month

Bit Torrent Changes to Use of UDP in P2P

RIPE Approves Private Market for IP v4 Numbers

A Discussion of the Operational and Economic implications of the Depletion of the Internet's Currency

Economics of IPv4 Numbers

International Bandwidth, Pricing and the FCC Deliberately Suppressing Demand?

A Note from the Editor on the March 2009 Format and Presentation

This issue has an Introductory essay by Mark Cooper and discussion; and article summarizing the state of play in emerging markets for IPv4 addresses; and eight weeks of symposium discussion - including a very long discussion of RIPE December 16 2008 address policy changes.

Please treat the IPv4 article beginning on page 15 as an introduction and executive summary for the long ipv4 discussion.

Text, URLs and Executive Summary: I have attempted to identify especially noteworthy text by means of boldface for **REALLY** good "stuff" . **Also the proper Executive Summary in this issue continues**. I hope you find it useful. Feedback welcomed. You will also find live URL links and page links in this issue.. (I am also no longer changing British spellings of things like fibre to the American fiber.)

Thanks to **Sara Wedeman -** see sarasworld.blogspot.com/behavioraleconomics/ for assistance with the masthead logo. Captain Cook now charts direction by looking at a compass rosette.

Coming in the April 2009 issue - **out about Feb 28 an interview with Harvey Newman will probably occupy at least the next two issues.**

I am omitting the contributors' page since a cumulative list may now be found at http://www.cookreport.com/index.php?option=com_content&view=article&id=121&Itemid=74