



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



© Cook Network Consultants

The Rise of Utility Computing in the Context of the Global Grid Could it Ever Offer Carriers an Opportunity to Cut the Umbilical to Voice?

This issue examines how increases in processing power and decreases in the cost of storage combined with affordable bandwidth on a global scale are enabling what is variously referred to as grid, utility, and cloud computing.

While users still have their edge-based infrastructure and the client server model of the 1990s has, by no means disappeared, **the maturation of grid computing into utility computing offers companies of all sizes the opportunity to use the internet as a more powerful and less expensive alternative** to client server computing. In many ways with these events, the network is becoming the computer. User's machines are plugging, like intelligent terminals, into services offered by Amazon and Google and increasingly others that run

on top of utility computing infrastructure and offer entrepreneurs ways to build services in software without having to buy and operate complex expensive hardware.

These changes underway at the level of the internet's "engine rooms" have profound implications. Are we, for example, looking at the emergence of a global computing utility that is becoming as important as the electricity grid? Also does the emergence of utility computing offer the carriers an inducement to transform their networks into data networks the primary purpose of which is to offer enterprises access to computing services of all types that could be safely and securely performed within the carrier's cloud rather than in house by IT staff at each separate enterprise. In such a scenario voice is no longer the major

Volume XVII No. 1, April 2008
ISSN 1071 - 6327

carrier product but rather a feature of all the data applications that utility computing enables. We shall see that there is some evidence that this the business strategy behind BT's 21CN, Global, IP Network's using 3Tera's AppLogic Utility Computing OS to run its own applications and then to offer services to BT customers.

First Some Background

Nicholas Carr, in his new book *The Big Switch: Rewir-*

On the Inside

The Global Computing Grid

Contents p. 49

Please read Explanatory Note
page 52

ing the World from Edison to Google, describes the evolution of technology in the post *Does IT Matter?* world. In this slender book Carr shows how the evolution of Google style utility computing is now enabling both enterprises and ordinary citizens to rely on second party providers of utility computing to do far more effectively and cheaply what hugely expensive corporate IT departments had monopolized before coming of Web 2.0 and high speed fiber links had made the location of computing resources less relevant.

Carr offers a useful metaphor for understanding what has happened by inviting us to look at the development of electricity and its impact on industrial development via the economy of scale offered via the electric grid. Power for machinery, whether water or steam or early electricity, had, like the services of a mainframe computer, to be at the location where it was used in the absence of effective transmission capability. However changes in technology including alternating current, enabled electric generation to move increasingly farther away from its customers. Consequently more customers could be served by a single generating plant. The result was a drop in price that encouraged further use.

The establishment of the first electric grid by Commonwealth Edison in Chicago enabled an even more radical drop in cost which ensured widespread use with the result that electric generation was never thereafter done in house. With both factory and home use enabled, society was transformed.

Since the bursting of the tech bubble, Carr shows how both the application of open source, and commodity priced hardware that becomes ever more powerful, has combined with the high speed optical links of the global internet to render computing as cheap and ubiquitous as the electric grid.

“What the fiber optic internet does for computing is exactly what the alternating current network did for electricity: it makes the location of the equipment unimportant to the user. But it does more than that. Because the internet has been designed to accommodate any kind of computer and any type of digital information, it also plays the role of Insull’s rotary converter: it allows disparate and formerly incompatible machines to operate together as a single system. . . . By providing a universal medium for data transmission and translation, the net is spurring the creation of centralized computing plants that

can serve thousands or millions of customers simultaneously. What companies used to have no choice but to supply for themselves, they can now purchase as a service for a simple fee.” p. 60

Moore’s law gives increasingly powerful microprocessors. These processors enable very cost effective virtualization – a process that also impacts computing architectures and drives them towards utility use. Carr explains: “ Because running a virtual machine is no different from running a software application, it consumes a good deal of the microprocessor’s power. until recently that limited virtualization’s usefulness. Running just one or two virtual machines would slow a computer to a crawl. – there’d be no processing power left over to do anything with the machines. But ordinary microprocessors have [now] become so powerful, that they can run many virtual machines simultaneously while still having plenty of processing power in reserve to run sophisticated business applications on each of those machines.”

“Virtualization breaks the lock between hardware and software that made client server computing so inefficient and complicated. A company no longer has to dedicate a powerful server to just

one application. It can run many applications on one computer, and it can even automatically shift the machine's capacity between the applications as demand for them fluctuates. Virtualization allows companies -- or the utilities that serve them -- to regain the high capacity utilization that characterized the mainframe age while gaining even more flexibility than they had during the PC age. It offers the best of both worlds." p. 76

"Virtualized systems that are shared by many companies are often referred to by computer professionals as multi-tenant systems. The name suggests a metaphor that gets at an essential difference between the client-server and utility models of computing. When you install a new system in the client server model, you have to build the equivalent of say a four-story building. But your building ends up being occupied by just a single tenant. Most of the space is wasted. With the use of virtualization in the utility model, that building can be divided into apartments and rented out to dozens of tenants. Each tenant can do whatever it likes inside the walls of its own apartment -- but all of them share the building's physical infrastructure -- and all of them enjoy the resulting savings." (p. 77)

On the very next two pages Carr discusses 3Tera as the prime example of virtualization in utility computing. Carr's book helps one connect the dot's -- especially one who has just dug deeply into 3Tera with a lengthy interview that follows this introduction.

Connecting the Dots or Why JP Rangaswami is Smiling Tonight

Carr points out that the way we access information is undergoing a new kind of change from a client server model to a grid model -- a change that is for the most part hidden from the ordinary user of the Internet. While we are still fixated on the fate of the telephone company, one phone company stands out for its forward thinking. BT seems to be ahead of the curve and I have been looking more closely at them than at others. For several years BT has been unbundled at the local loop and free to think about next generation non voice services. I wrote six months ago how JP Rangaswami and Jeremy Ruston have open sourced BT's web services to voice and data developers.

But JP was brought in from Dresdner as Global CIO of the BT 21CN IP network -- a title that sounds to me as going a

bit beyond the objective of just making a new platform for mobile voice services.

We have seen over the past 18 months that BT has acquired many firms on the Eurasian continent that are both involved in fiber networks and enterprise services. I noted 15 months ago that JP's old boss at Dresdner and new boss at BT Al Noor Ramji was quoted in June 2006 at a telecom conference in Spain as saying that he regarded Google as BT's major competitor. At the time that very likely seemed to some to be a daft statement. What does a search engine company have to do with a phone company?

It seems that Carr's *Big Switch* may have the answer. **Google not as search engine but Google as grid for utility and cloud computing.** By the summer of 2006 BT was likely already talking with 3Tera. In any case on June 6 2007 *Network World* announced a multi year contract with "3Tera to deploy its AppLogic grid operating system to support future software-as-a-service and other offerings.

"BT, . . . is deploying the AppLogic system internally, **but it will be used for future services such as software-as-a-service, says Venkat Raju, senior strategy consultant at BT.**

He says BT is also conducting customer trials using the AppLogic system.”

“While other service providers have been offering utility computing services for the past three to four years, those offerings are generally manual intensive and do not allow customers to directly request more bandwidth or to turn on an application themselves. Instead requests for more bandwidth or for a new feature are directed to the carrier, which then fulfills the request.”

“3Tera’s AppLogic system automates these processes.”

Editor: see Carr’s remarks above on virtualization – and 3Teras that follow.

“BT conducted a detailed search for a platform that would work well with its 21st Century Network initiative, Raju says. “3Tera’s product came the closest to meeting our requirements...it specifically has a virtual modeling environment that was very important to BT.””
<http://www.networkworld.com/news/2007/061207-bt.htm>
 |

BT, by effectively becoming a part of the global grid through its own extensive fiber and its ability to let any enterprise make encrypted connections to data centers running AppLogic, is setting

itself up to offer Utility based IT services to global enterprises. As readers of the 3Tera interview that follows will see **AppLogic, as an OS is specifically constructed to make it possible to move generic enterprise applications off corporate servers and on to utility grids with no changes in code.** BT is seems is beating all the other national carriers from the model of telco as private pipeline for the voice and data traffic belonging to others to carrier as purveyor of commodity priced, utility, grid-based software and virtualized services.

Utility Based IT Services to Global Enterprises – the New Carrier Business Model?

A skeptic might say that enterprise IT types will be slow to embrace radical new changes in their operations. Indeed – except that BT’s timing on this is impeccable. JP – I am sure – remembers well what happened in 2001 after Al Noor had left Dresdner and he was promoted to Al Noor’s position as Global CIO. As the bubble burst his budgets were slashed. IT must do more with less. JP did. Blogs, wikis, and a major embrace of open source software. The proprietary and expensive was jettisoned with great rapidity. Soon he

was doing more with a fraction of the cost and the staff.

Anyone who reads the financial news can see that enterprises in general and banks in particular are now standing on the edge of another financial abyss. Their IT budgets in 2008 will surely be slashed. I’d wager that a great many will try BT’s AppLogic services as life preservers.

No wonder JP Rangaswami is smiling tonight. Smiling not I am sure at the discomfort of other IT managers, but at his good fortune at being in a position to push the barriers of technology in a way more creative than most of his counterparts by positioning that technology to help enterprises use the brain power of their employees in the most optimal way.

On February 9th Al Noor Ramji asked: “Any comments from people who would like to argue whether or not Google is in this space?” My response is that I see no evidence yet. Google has publicly talked up cloud computing in the research and education space. For the moment it seems that Google publicly at least is not. Could they jump into this arena? With great ease. Are they interested? They would have to be stupid not to be. They are definitely *not* stupid. The last segment of this month’s

discussion focuses on their moves to increase the penetration of Google Aps into enterprises. Google would be foolish not to have something like this in mind.

Meanwhile, on February 23, there came another clue with regard to BT's timing. Josh Snowhorn, VP at Terremark gave a presentation on Data Center Power trends at the 42 Nanog meeting.

<http://www.nanog.org/mtg-0802/presentations/Snowhorn-Power.pdf>

While the whole presentation is very informative, the first slide is quite an advertisement for BT's approach to Utility computing.

"By 2008, 50% of today's Data Centers will have insufficient power and cooling. By 2009, energy costs will become the 2nd highest cost of a Data Center - source: Gartner **By 2010, half of all (Enterprise) Data Centers will have to relocate or outsource applications to another facility. - Source Data Center Institute.** During the next 5 years, 90% of all companies will experience some kind of power disruption. **In that same period one in four companies will experience a significant business disruption.**"

BT, it would seem, is riding several positive curves in its utility computing strategy,

Meanwhile Back at Home

Finally in explaining what the evolution of the global grid of utility computing offers, Carr does a good job of equating access to this grid as being as important as access to electricity a century ago.

What he doesn't do however is adequately distinguish the critical distinction of broadband always on access from dial up. This is the only way that one can have the benefits of the computational grid.

What does become clear is equivalence of access to computing now as to access to electricity a century ago. A modern viable place in any economy without it is not possible. His extensive catalogue of what can be done is useful to the policy advocate in building an extensive portfolio of the reasons why anyone without access is economically disenfranchised. I would argue that the similarities to the electric grid and computing grid are so extensive that the two together make a better teaching tool than Mike Bookey's highway analogy of 2006. I am hopeful that the

lesson to an Obama administration will be as obvious in 2009 as access to electricity was to President Roosevelt in 1932.

The Network Becomes the Computer: Grid, Cloud, and Utility Computing

An Interview with 3Tera Traces the Origins Functions and Benefits of an OS for Utility Computing

Editor's Note: I interviewed Barry Lynn, Chairman and CEO of 3Tera and Bert Armijo, Senior VP Sales and Product Management of 3Tera on January 22, 2008. (Their bios are found at <http://www.3tera.com/management.html>).

Grid, Cloud and Utility Computing

COOK Report: Let's start by getting your point of view of grid, cloud and utility computing – their origin and development and the differences between them. How do we get from grids to where you are?

Armijo: Having been a passing idea in the sixties, utility computing re-emerged in about 2000. My first acquaintance was at an IBM Infiniband conference in 2003 and I remember the conversation that emerged in the wake of Nicolas Carr's article "Does IT Really Matter" in the aftermath of the dot com bust.

COOK Report: Which came first? Grid or utility computing?

Lynn: Grid has been around for a long time in its use by the scientific community for calculation-intensive problems.

Armijo: But the users of the grid did in fact have utility computing. When the concept of utility computing caught on as a buzz word five or six years ago, everyone wanted to offer it and there were all kinds of companies attempting to do so, but the technology needed to address the problem wasn't ready yet.

COOK Report: How do you define Utility Computing?

Armijo: It is the ability to access and control remote computing resources that can be used for general-purpose application. You need to be able to run anything from SAP, to Oracle, to Apache, to number crunching. The issue with the grid, up until the time 3Tera came along, is

that it required reprogramming. Consequently it was only really tailored to technical computing and related number crunching applications. You needed to write your code specifically for a grid OS that ran on whatever grid you were going to subscribe to. You could submit your job that would be scheduled and then run.

The computational resources were vast but the bandwidth connecting them had fairly high latency. It did not lend itself to transaction processing. It was more compatible with parallelized scientific calculations. Secondly, your job was expected to end. You could not steal the grid forever. Other people had computations to run as well. So, as long as your job fit these characteristics, you had pretty good utility computing. If you had a program that required modeling, you didn't need to build your own grid. You simply submitted it to whatever grid your school subscribed to.

The general-purpose world was not this fortunate. They still had to build out resources every time they wanted to deploy an application. As an industry, we had started in 1995 -96 to move away from large SMP systems (Symmetric MultiProcessing, that use multiple CPUs within one box) like Sun E10000 servers to commodity arrays and to distributed computing. When we first started doing this, the distributed systems were very small and the big system administrators could run them almost as easily as they could run the big SMP systems. The cost delta was enormous. A million dollar SMP system versus 10,000 dollars worth of servers. The return on investment calculation for making that move was pretty simple.

But in the process we lost something without realizing it. We lost the structure of applications. When you ran on a big SMP system, you developed your code on a Sun workstation. You divided your application up into little pieces called daemons. You had a shell script that would launch these daemons; connect them together through sockets, define their IO sources, and would control their life cycles. If there was an error in the daemon, the shell script would handle it.

It all worked very well. When I wanted to deploy, I simply tarred the directory and copied it over to the larger SMP system in the data center and it ran. I did not have to go buy hardware or do anything else in order to deploy. The infrastructure of my application existed within this shell script that ran on the SMP system. To Sun's credit when the move toward commodity servers started, they were standing up and saying: "you don't understand. You are leaving something behind". But no one listened.

Now, here we are more than ten years later and these distributed systems have become enormous. Instead of a couple of dozen servers, it is now routine to see people with hundreds or even thousands of servers connected together and performing a single application.

Lynn: I think Bert gave you a great 500 foot view of what's going on now. Here's a 50,000 foot view. My **definition of utility computing, would be an on demand service like electricity or water. What you want to think of is packaging a bunch of CPUs and storage as almost a metered service.** There is a low initial cost to acquire it. You don't have to buy hardware because you are basically renting it. If you have sudden peaks in demand, you don't

have a delay in having to acquire physical machines. Instead you can just dynamically allocate the resources you need.

Armijo: Grid computing is a distributed computing technology that you most often see in many utility computing models. You also see virtualization as a technology within most utility computing models. To me the definition of utility computing is an **on demand, metered service that can dynamically go up and down without end users having to physically acquire their own hardware. It also has to be able to run standard general purpose software.** Otherwise all you would be doing is shifting your costs from building out hardware to writing specialized code.

To me grid computing was always been a form of distributed computing with nodes scattered across domains. Utility computing is basically the bundling of those resources for sale to the public.

COOK Report: And what might those applications be?

Armijo: The most obvious are search engines. They tend to be the largest ones out there. Video services are another. We have a customer getting started right now that is a video training distributor.

This customer will probably take 20 or 30 servers, but the storage requirements are astronomical. We have had customers looking at 100 terabytes, 500 terabytes, and even a petabyte.

COOK Report: I note that BT is a customer of yours. Would you comment on that?

Lynn: The Web 21C project was the initial one that they brought us in for. They are now using us mostly for internal applications. Their intent is to use 3Tera and AppLogic for both internal and hosted applications. They need to be able to say that they "eat their own dog food," before they offer it to others.

Cloud Computing?

COOK Report: So where does cloud computing come in? I say this jokingly but wonder if it is a Google conspiracy to confuse people?

Armijo: Cloud computing is a fairly new buzz word, and as always with these new technologies, rather nebulous. When we talk to customers about cloud versus utility we see some trends. Different trends among completely different customer bases.

If you are considering cloud computing you are willing to have someone

else control your resources, their location and how they are scaled. Cloud computing is about being willing to give up control. A good example of this would be people who are running on Salesforce.com. They don't care how Salesforce runs or how it scales. What they care about is that they are getting their SLA and that it is available to them when they want it.

On the other hand, utility computing is precisely the opposite. **Utility computing users are control freaks because they are building services; they need to have control of their resources. Thinking about it this way, would lead us to say that utility computing is the appropriate platform for building clouds.** If you are going to build the next Salesforce, you are going to have to know about latency and what is running in which data center. You need to know that latencies are consistent and you need also to be able to schedule these backups and know they are taken care of. **You need to have very precise control of the infrastructure necessary to build your cloud so that your users can trust you when they give up control of their data to you.**

For example, people talk about Google as running a

cloud. The question then becomes is there any way that you could build Salesforce to run on Google? Absolutely impossible. Could you build Facebook on Salesforce? No. It is not designed for general purpose computing.

On the other hand, with a utility computing service, no matter whether you talk about 3Tera or about Amazon, you could actually build Facebook or Salesforce on either platform. This is where we see the delineation coming with the two kinds of services.

More and more, when we see people writing about cloud computing they are writing more about consumers than businesses. You look at people writing about the MacBook Air as the definitive enabler of cloud computing. Well, this is a laptop. What they are trying to say is that it is about consumers storing all their data online, using some software service as opposed to keep it on their laptop.

In that respect, the definition for cloud computing is becoming consumer centric. Still, it is really about giving up control – about not having to know where all these things are – who is running them or where they are running, and for consumers this may make perfect economic sense. Utility computing us-

ers want to know as much as possible about how their infrastructure runs – even though they don't own it.

COOK Report: Utility computing then is an environment that you can use to build a special purpose application.

Armijo: Yes. As a consumer I love not having to worry about the precise structure of the services I am using. Also I need to be able to tailor and integrate them into my business as much as possible. The vast majority of what I do as a consumer is already in the cloud and has been for years. But as a business owner, I want to know how my systems are running, where they are running and so on. Also, I need to be able to tailor and integrate them into my business as much as possible.

Lynn: And if I go back to my 50,000-foot view of things: People are touting cloud computing as the ability to simplify IT management.

COOK Report: With BT would it be fair to say that they are using your utility computing offering to build various network applications that others will use as a cloud?

Lynn: I would say what they are using it to turn hundreds

of services that they offer now in their data centers into utilities for their own internal use. They want to be able to have an internal utility computing offering for their own IT staff. When they start moving it to their own hosting business, they will then be offering utility computing in the form of virtual data centers and virtual servers to their customers. I am sure they will also have a cloud computing service where people can request services and have applications with infrastructure that is invisible to them and be able to migrate those applications geographically.

Difference Between Amazon EC and 3Tera

COOK Report: Could we move on to Amazon? A year ago I met and interviewed Kevin Nethercott of Lignup who built a Call center service to run on Amazon's Elastic Compute Cloud. Help me understand the difference between you and Amazon.

Armijo: EC2 came out as a Beta in August 2006, but Amazon web services for creating virtual storefronts have been around since 2002. EC2 gives their customers the capability of deploying virtual server instances on Amazon's infrastructure. They refer to virtual machine instances as virtual private servers and

sell them in different sizes as a metered service.

Amazon Web services started as a way to get retailers into Amazon and to get them to use Amazon as a platform for selling on-line. It started with the Toys-R-Us deal. Toys-R-Us had tried twice to get an online presence and had failed both times. Amazon was able to make them very successful. This was the beginning of Amazon web services. After this it was about what services are necessary for people to make it easy to bring their electronic stores out through Amazon instead of trying to build their own. It has become an enormous revenue stream for them.

COOK Report: But Lignup was using them to put a virtual call center on line. I have the impression that they are moving with EC2 away from "just" web services.

Armijo: They are. When Amazon launched EC2 they were honest in saying that they had built out a platform to run Amazon and they were going to allow people to run on the spare resources that Amazon has in their data centers. Having built EC2 for their own needs, the system did not need to conform to expectations of how standard PC servers run. One of the most obvious examples of this difference is the lack of

perpetual storage for virtual servers. As a result, when you use an EC2 virtual server and that virtual server image stops all data on the disk drive associated with it is lost. It simply goes away. Since this isn't the expected operation of a standard server, most software isn't written to operate in this way. As a result, users have to find a way for their software to write the data to another service, such as Amazon's S3.

If you're launching a service that people can use to leverage your existing infrastructure investment, that type of difference is OK. People will either use it or they won't. **But if you are starting from scratch to build a commercial utility computing system and you are going after the mainstream market, you better be able to run standard software.** We take this as a given, and this is a principle difference between Amazon and where 3Tera comes from.

3Tera built AppLogic from day one as a commercial utility computing platform. Therefore we took it as a basic assumption that everyone must be able to run their standard software without any changes. Frankly, it is incredibly difficult to get a utility computing system to run standard software, and this is particularly true when it

comes to storage. Our success revolves in part around the fact that our engineers have built storage systems in the past.

Amazon took a different approach. They clearly decided that for their purposes they did not need to run standard software and consequently stepped the storage problem by building a storage service instead.

If you are building an application from scratch as a developer – lets say you are putting your call center together and you don't mind writing your software so that it will specifically only run on Amazon, then you can do that. However, you also need to recognize that if you decide you don't like Amazon or if Amazon closes their service, you are in a bit of a bind, because you have written your software specifically around Amazon's interfaces and you'll have to rewrite it before you can move somewhere else.

3Tera introduces no new APIs, so you don't have to rewrite code in anyway to run on AppLogic. Therefore, you can always move your applications to physical servers and it will run there in the same way.

Genesis of 3Tera

COOK Report: How did 3Tera get started?

Armijo: There is the main engineering group at 3Tera. Barry and I are not part of that. These guys have worked together for a long time in many companies. The last one was called Z-Force which is now called Attune Systems. There they built the first file storage switch. Unlike traditional storage switches (which operate at a block level) the file switch directs requests based on files. For example, instead of mirroring and striping blocks of a whole disk, you could choose to mirror small image files but stripe large streaming media files. In addition, the file switch aggregated multiple file systems and presented them as a single file system, with the combined capacity and performance of both systems. This is a type of virtualization, called aggregation – it allows users to see the combined resources as if they are one big pool, rather than hundreds of discrete pools. Using aggregation simplifies management and operations of even the largest systems. In 2002 they demonstrated the world's largest and fastest commercial file system, built using the file switch and hundreds small Dell NAS appliances.

Lynn: In other words this was managing files rather than the physical storage of blocks of data. They actually invented the file switch. Before this, the way you managed storage was to manage blocks of data on the disk.

Armijo: Meanwhile, I had a company in Northern California called Topspin. We were building what we considered to be the world's first data center network virtualization switch. Topspin is now part of Cisco which has kept the name of the product VFrame. The idea was very simple – any server in the data center should be able to take the place of any other server in the data center.

We had a switch that controlled the storage system and the networking for all of the servers in a data center. If a server died, we would simply choose a spare server and map the identity of the new spare server so that it took over the activity of the old failed server. We did not use virtualization at all. We did this by mapping all of the network and file and storage items. This was very valuable for people who were running large data centers organized around a specific task because it kept the system up and running. For example our largest customer was Merrill Lynch for their currency trading floor. This is a huge number crunching ap-

plication where the time required to bring up a spare server is incredibly important because it shaves the time necessary to make a trade. Time for them is really money.

COOK Report: Did this capability become a component of what later becomes **AppLogic**?

Armijo: As capability, yes. Although what AppLogic does is much more and the technology used to achieve it is totally different.

To continue the story, after I left Topspin and met up with Peter Nikolov and Vlad Miloushev who were just forming 3Tera. We decided that we both wanted to work on the same problem - how do you scale online applications? We had seen many customers trying to scale on line applications and having enormous difficulty doing so.

We played around with some ideas and at some point Vlad exposed to me the technology 3Tera has built and I recognized how it applies to the drawings I had been sketching at Topspin but couldn't figure out how to implement. This was the genesis of **AppLogic**. Here was someone who actually had a concept for doing what we had been trying to achieve at Topspin. This was over Christmas of 2004.

In February 2006 AppLogic went into beta and for the first time we introduced users to the system. It was April of 2006 that the first customer application went live online. That customer *International News Media* is still with us. Today, AppLogic is at version 2.3. We have been steadily bringing out upgrades and new functionality, completing the system vision. We have things on our road map that will take us out through the next three years in order to implement the vision that we have outlined.

This is a brief history but the concept was not so much let's go build a grid operating system but rather figuring out how we help people build online applications.

From Scaling Online Applications to a Grid Operating System

COOK Report: But who was doing what?

Armijo: Before 3Tera Vlad was doing the file system management and I was doing the data center virtualization.

Lynn: **Vlad founded 3Tera as a company designed specifically to scale applications.** As Bert says he was trying to do the same thing at Topspin but his approach was different.

So I think it is fair to say that when you and Vlad hooked up, you Bert had the knowledge of data center virtualization and scaling on line applications and Vlad was already developing stuff for doing that scaling.

Armijo: You are right. **It wasn't that we started out to build utility computing or a grid operating system but that we started out to solve the problem of how to scale online applications. As we started going forward the pieces started to fall into place about how these things could and should operate and where the most difficult and time consuming problems were that people were trying to solve.**

It took us a year to admit that what we were building was an operating system. You start these things in a piecemeal way. Your users need to be able to do 'x' and use 'y'. They need to build on top of inexpensive hardware and then, at some point, you simply realize that what you are building is a grid operating system. When you realize that it is very empowering because you can then think in terms of how you actually complete the system.

When we finally announced AppLogic and called it a grid operating system, a lot of people didn't accept this. They asked where our interface was and where were the APIs to which people could write? People also complained that if it was a grid, it had to be designed for number crunching, which, of course, it is not. So, we had some initial push back.

No one had ever used virtualization to build a grid system before we announced AppLogic. Looking back it is a lot of fun to see us having this impact.

COOK Report: So, if you begin to think about not just scaling systems but building an operating system for use in scaling systems you have a new conceptual framework that allows you to progress in more productive and creative ways?

Armijo: Yes. We focused on the problem how to scale online applications for a year, before actually admitted to ourselves that we were building an operating system. At that point we had enough of an understanding of how the problem was affecting the user and how it could be solved. Saying it was an operating system actually made things obvious and enabled us to complete the product properly, adding the needed finishing touches.

There are a lot of products out there that focus on pieces of the problem. For example, image managers launch images. They come at this from the point of view that we get started by deploying things. But they lack a holistic view of how the system should function.

COOK Report: So you began to look at a basket full of functions from the way in which they were interrelated, and from I suppose a structural point of view?

Armijo: Absolutely. **AppLogic is the only product out there that understands the structure of the applications it runs.** Because of that we are able to do things that other people look at and say: "someday we will do that too". For example, with AppLogic you have the ability to move applications between data centers. This is an amazing thing when you realize what we are talking about here. We have a customer who builds a search engine and they were able to send us a copy of their application. In other words, sending us a search engine with all of its data, its load balancers and its networking infrastructure – complete, packaged and ready to run; so, that we could help them debug a particular problem. An application like this is not a small, it

normally runs on a few hundred CPUs.

Today, to move an application of this size and complexity is a completely impossible task for anyone else except the AppLogic users. That's because when we started our goal was to help people scale applications instead of just to help them run software.

VMware has a specification called OVF (Open Virtualization Format) that they pushed through the DMTF. One of the things that they claim they want people to be able to do is to help them migrate a virtual machine. We can already migrate entire virtual systems comprised of multiple virtual machines from any AppLogic equipped data center to any other AppLogic equipped data center.

Lynn: I want to stress again, one of the key things that Bert said when he was giving you the example of how we could debug an entire application, since we had access to its entire infrastructure: What we create with AppLogic is the ability to have an image of an application and its entire infrastructure. Whether you are running that application on one server or 100 servers makes no difference. It is merely a matter of allocating the resources you need for the load that you are trying to process. Therefore migrating an

application from one data center to another will include an image of the entire infrastructure it runs on. If you have a typical n-tier online application, it has load balancers, firewalls, app servers, web servers, database servers, etc., and all the data in the database running on your AppLogic grid. With a single command, you can migrate that to another data center from where you left off with an image of that application and all its infrastructure intact.

The Infrastructure

COOK Report: OK. You have a utility computing operating system that is developed and runs on commodity servers and does good things, but who owns all this? How do you acquire and pay for this? You mention AppLogic data centers. Where are they, how do they come into existence?

Lynn: We don't own data centers or servers. AppLogic is currently installed in 17 data centers in North America, Europe and Asia. There are three ways to become an AppLogic user. 1) We have a couple of customers like BT who run AppLogic on their servers in their own data center. 2) Our core customers are hosting providers. Their main business is to run, operate

data centers and rent servers. We work with hosting providers who co-brand and resell the utility computing service using AppLogic. 3) We also have handful of customers who sign up directly with. They still use the service from one of our hosting partners. The only difference here is that we are the first point of sale and the first point of support.

COOK Report: So if I am Equinix, or the Amsterdam Internet Exchange or one of the other big internet exchanges these exchanges could in effect "host" AppLogic Data centers?

Armijo: Yes. AppLogic does not require any unique of specialized infrastructure. It is simply Intel or AMD servers, direct attached storage and gigabit Ethernet. There is no unique hardware whatsoever involved in building an AppLogic grid. We are trying to enable operators – people who already own large numbers of servers to add utility computing services to their product mix.

I think the extent to which we might be embraced by a large internet exchange like AMSix or the Equinix exchanges is the extent to which these exchanges want to get into the business of providing utility service in addition to serving as large scale switching fabrics.

One of the reasons that we are excited to work with BT is that as utility computing moves out of the web 2.0 start-up environments and into the more general business environment this is the customer base being served by BT and other telcos. These companies are already providing the data communications needs of every enterprise from Reuters right on down to the local liquor store. And as a result of this, they are a natural ally for us in providing utility computing services to their broad customer base. They have the skill set in running the operations. They have the skill set in terms of providing the support and they have the skill set in terms of making the sale.

COOK Report: And they would have the physical environment as well that is necessary to support your software?

Armijo: Absolutely. But it is not just supporting the software; it is also supporting the customer who is using the software. This is why we were excited to sign up with BT because we think that as a major operator who is dealing with that customer base already, they can address that customer base very easily as they start to move util-

ity computing into it.

COOK Report: For an enterprise using your services what are the risk management issues? Is there any question to worry about of putting all their eggs in one basket?

Armijo: With AppLogic there is not just one grid. You can have many grids. For us a grid is an organization of resources and not a universal thing. While we are running in seventeen data centers right now it is quite possible for users to have presence in multiple physical data centers with AppLogic and, as I mentioned, you can move applications back and forth almost at will. We do have customers who are using AppLogic for disaster recovery scenarios or as we like to call it business continuity. They run at a second data center and do data base replication between the centers, so that should one center fall off the face of the earth, as has happened with disasters like hurricane Katrina, they in fact will have a copy of their application up and running and transparently accessible to their users.

The second data center could be running in their own facility or on a different continent. As Barry said, we do license AppLogic to end users as well. Let's say you are a bank. The ability to run the

same platform in your data center and to get that platform hosted in more than a dozen other data centers in other parts of the country or world is quite attractive. It gives you the opportunity for flexibility in your business decisions going forward.

Do I want my payroll to run in my data center? I may at the moment. But let's say I have some large increase in the volume of my other jobs and need to burst out of my center. Suddenly, I can take my payroll running on AppLogic and move it elsewhere with a single command. Moreover, I can select exactly which center - the cheapest one or the one that offers me certifications; the one that has special security or the one that is HIPPA compliant. By partnering with many different hosting providers and Telcos, many additional types of services become available. These services can all be layered on top of the utility offering by the vendor. This gives our customers valuable choices in how they will apply their resources.

The Business Model

COOK Report: To begin to wrap up, what is your business model at this point? Where are you going?

Lynn: As I mentioned before, we have a few categories of

customers and sources of revenue. There are hosting providers who co-brand and resell our product. They can resell anything ranging from virtual private servers to virtual private data centers to subscription to a service stack on which users build their own online applications.

We have small and medium business customers who sign up for the service directly with us. We direct them to hosting partners but we provide first line of support for them.

Most recently, we started offering AppLogic for the enterprise. This is our software with an enterprise-licensing model attached to it. In this case instead of installing it on a metered basis we install it on a server-licensed basis.

This is the business model today. I think the future is utility computing and cloud computing and obviously the ultimate goal for us would be to become the de facto standard for utility computing. I think it is going to be a while before there is a standard for utility computing. However, we are the only people out there that provide a platform that is agnostic to the operating system, the programming language, GUI, database, and app server – meaning that you can run anything on it. There is nothing proprietary about it at all. Consequently,

the end game for us is to be seen as the major player in utility computing in the coming years.

Armijo: There are two parts to this. One is the answer to what we are trying to do – we want to build global network of partners who are running utility computing services. Whether they are businesses or banks, insurance companies or government agencies, doesn't matter to both of us.

Second part is looking forward. Our view of the world looking ten or fifteen years from now is that it makes absolutely no sense for most businesses to own a server. It just doesn't add any value to their business anymore. The notion that we have to have servers in our business to run software is archaic and is going away.

COOK Report: **And the reason that this makes no sense is that if the business has the Internet, to which to connect the functionality, the server is easily achievable by other means?**

Lynn: Yes, for two reasons. **First, if you can connect and get the right amount of IT services on demand, then it will be much more cost effective to do so than owning and main-**

taining the services yourself. But, much more importantly, you can greatly decrease your time to market, thus, accelerating revenue while increasing profit margin. Secondly, why do you have to manage them? For the most part, doing this is not the core competence of any business.

I talk to people all the time about why businesses shouldn't own data centers and servers and they look at me like I am crazy because, of course, most businesses do own servers and data centers. So the analogy that I use is this: If your company business is to sell things online – do you buy trucks and airplanes to deliver them or do you use FedEx and UPS? So if you are in the health care business or manufacturing or supply chain business, why does it necessarily make sense for you to own and operate data centers?

COOK Report: Where do security issues come into this?

Lynn: If you own your own data centers, you are shipping stuff from point A to point B, so you will have, in fact, a bigger security issue than you would if were just using a service. I think there are a lot of people out there who try to resist what we offer by saying I can't lose con-

trol of my own data and data center because of security and privacy issues. The reality is that the more you try to hold on to your own data centers and services the more privacy and security issues you will have. In addition to hackers knowing who you are and where your operations are, you have the additional concern of maintaining and replacing your physical data center infrastructure periodically. Since this is not your core competency, you will likely be tempted to stretch out your investment.

The same issue is not true in running in these hosted data centers. If you are in the health services business, you cannot afford to be continually upgrading the physical and logical security of your data center. Better stuff comes out every year and you won't be able to afford one year replacement cycles. If you are in the computing business, though, you can't afford not to do this. I argue that by giving up the servers and data centers to the people who do nothing but this for their living, is actually much more secure than if businesses did things on their own. This way they have access to the best physical and logical security techniques in the business.

COOK Report: What about encryption?

Armijo: If you are moving your application from one data center to another, it is absolutely encrypted and signed with a public and private key pair. It is as secure as any communication can be that is running over public lines. If you wanted anything better you'd have to have private lines. Some of our customers do. We are talking with a couple of banks that, I guarantee you, have private lines between their data centers. In terms of adding it to an individual application, yes that is just a drag and drop function.

Now, an additional difference between us and Amazon (since they are the other utility computing system that people seem to understand) is that with Amazon when you use them, you are, by definition, running on shared hardware. You can never have an entire sever to yourself and you can never find out what else is running on your server. All of your storage is by definition also shared. Anything written to a disk is written to a shared disk.

With AppLogic we have taken a different approach. When you sign up with AppLogic, a grid is created just for you. It is your grid. The resources that you use are yours alone. No one else will be storing data on any disk to which you

have access. No one else will be running any kind of applications or code on any server that you are using. The servers that you are running for the time that you are using them are completely yours. Only your data will pass across those backplanes.

Now, let's say you have a traffic spike that caused you to select 100 servers instead of your normal fifty. When the traffic goes back to normal, the extra fifty servers allocated to you disappear back into the "ether" to be assigned to other people's grids. We wipe them clean automatically when you are done, but while they are on your grid, they are yours. No one else is going to touch them.

So, from a security point of view, if you need to be able to pass an audited test for running credit card processing, knowing that these are your own unique machines will enable you to do that. While if you are running on a completely shared infrastructure, like Amazon EC2 where you don't even know what server you are running on or who else is there, is not the direction you want to go in.

As a result of our approach to security we are able to guarantee performance. On an AppLogic grid if you launch your applications 100 times in a row, you will get the

same performance 100 times in a row. If you are on a shared infrastructure, you app goes to one server where there is no one else running you get all the spare cycles and wonderful performance. The next launch goes to a busier server and your performance changes. Under these conditions it becomes very difficult to plan out any uniform quality of service level that your application will consistently reach, if you don't have control of all your

resources. Under AppLogic you can always see what is running on what server. During testing you can actually force the failure of servers so you can test the resilience of your application to hardware failure, so, that you can understand how it will operate under various load conditions. This is all about being in control for the operator so that he can reliably run large online systems as part of his everyday business. Consequently, the CIO can go to

the Board of his company and reliably vouch for the fact that the technology he is using will not blow-up and that the company can shift critical parts of its business to a utility computing model.

Symposium Discussion January 16 - February 14 2008

In the European Union - Regulation Will Persist in Order to Guarantee Open Access

Lars Hedberg: On January 16: EU-commissioner Viviane Reding gave a good speech in Amsterdam regarding issues of broadband in US and Europe.

COOK Report: Quoting Reding: "In my discussions last week, I was surprised how critical the industry is of broadband services in the US. I found a widely held view that the European regulatory framework and its emphasis on access obligations to open up competition is not at all the impediment to investment and innovation that some market players claim, but actually the main reason Europeans are ahead of the USA in terms of usage of new web services such as social networking sites and virtual worlds."

"The regulation of access networks keeps open the potential bottleneck to competition in broadband markets. European rules have provided a real choice for consumers."

"The European model is empirically proven to promote not just choice, competition and innovation but also in-

vestment: in 2006, investment in the EU telecom sector reached another peak of over €47 billion, 5% up on 2005. This was the fourth year on year increase since 2003. European investment is at least at the same level as other major regions (Asia Pacific: €44.3 billion and North America: €43.7 billion). [snip] **How we treat next generation access is therefore the single most important policy question in the telecoms sector today.** We have to create incentives for investment whilst making sure that no-one (and I insist on this no-one), can be in a position to foreclose the market."

"But there is an important new question to answer. As we look at next generation networks offering high bandwidth, available everywhere, cheap and open to innovation and competition. Where will the investment come from to move us from legacy copper networks to new high bandwidth networks?"

"It is clear that next generation core networks are interesting to market investors: we have seen investments in

such networks across most of Europe. But access networks are more difficult for two reasons. The cost of the upgrade is high (between 300€ and 1000€ per household on average, in many places even more!). Second, access is the bottleneck and therefore it creates competition problems. If one player can get hold of this bottleneck he will have the means and motivation to shut off the market to other players - or at least give access to who he wants and on the terms he wants." [snip]

"But there is an important new question to answer as we look at next generation networks offering high bandwidth, available everywhere, cheap and open to innovation and competition. Where will the investment come from to move us from legacy copper networks to new high bandwidth networks?"

"It is clear that next generation core networks are interesting to market investors: we have seen investments in such networks across most of Europe. But access networks are more difficult for two reasons. The cost of the upgrade

is high (between 300€ and 1000€ per household on average, in many places even more!). Second, access is the bottleneck and therefore it creates competition problems. If one player can get hold of this bottleneck he will have the means and motivation to shut off the market to other players - or at least give access to who he wants and on the terms he wants." [snip]

"We can see that the choice of which type of Next Generation Access is offered is not neutral. Market investors have a strong incentive to keep investment costs low and where possible increase their control over the access bottleneck. That is why in my opinion regulation will and must persist, so that open access is guaranteed through effective non-discrimination obligations. The remedies will however have to be adapted to the access solution that has been implemented."

January 17 **Jim Kayne** on Silicon Investor: (as imported by Frank Coluccio):

In 1986, the Greater Vancouver area was host to Expo '86. In preparation, the first municipal high-speed transport was created: SkyTrain (sort of an overhead subway). In the following 22 years, areas parallel to SkyTrain have enjoyed a tremendous renaissance, experienc-

ing both residential and commercial growth far in excess of median growth in other areas.

Fibre buildouts need not be accompanied by some kind of telecoms land grab (thus creating more incumbents) - but rather, the exercise of imagination and creativity in proposing solutions. Reding recognizes (as we all do) the need for ROI. **What nobody (except our Asian brethren) seems to grasp is the fact that telecoms is infrastructure subordinate to the national and public interest.** Like many others, she uses the term regulation as the ultimate arbiter, when it is not. Regulation is a subset of the national government's policy. Government has the power to exercise eminent domain.

<http://supreme.justia.com/constitution/amendment-05/19-national-eminent-domain-powder.html>

Not that such exercise should be punitive or exclusionary (again, as demonstrated by Korea and Japan): incumbents can be given the option of playing the game by the new rules, or not. But if the new rules are structured properly, they'll be quick to realize that it's play, or die.

Creation of transport and capacity is a self-fulfilling proposition; contrary to popular thinking, excess ca-

capacity is the desired end. The uses for that capacity will be defined (and will change) by the accompanying ecosystems, themselves. Reding appears to be justifying fatter pipes by the existence of unsatisfied demand for new services: that's only partial justification

It's not as simple as "build it and they will come". Rather, it's the ability to generate flow. Impeded flow (read scarcity) stifles growth. You can't tell what will develop; whether it will be replicative or novel; primary, secondary, or ternary; interactive locally, globally or in-between.

Hendrik Rood: I was at that event in Brussels, where Reding delivered the speech. Nico Bakken was there too.

Cross reading the paper version of the speech I observe she made some occasional additions with respect to mobile broadband, and also some announcements/thoughts on the additional policy letter the EU would release in August 2008 with respect to broadband / FTTH.

One of the prime points she made is that she would ask for a more detailed set of statistics (past and above current level, that tend to count mainly broadband access lines). Reding stated that the

EU wants to add statistics collection with regard to (from memory) - access lines - bit rates up- and down and prices. And some more on quality related issues (she did not elaborate on that. Perhaps over subscription factor might also be one).

Official statistics on the state of broadband in Europe vary

widely in quality and depth. In countries like the Netherlands legal requirements on low administrative burdens for corporations force our regulator to take a 'de minimis' approach. So the EU could kick that low level a bit up to give more insight.

Reding also made a side remark that the EU would

probably set a broadband penetration goal of 30% lines per capita for those countries still far behind the leaders. E.g. Greece, but most southern European countries are trailing.

Does the Credit Crunch Exacerbate American and European Tendency Toward High Return Short Term Investment?

Monetary Resource Allocation in the Credit Crunch

Robert Cromwell: I heard on Monday from a very creditable source that the sub-prime crisis is negatively impacting allocation of money for funding utility capital projects within the larger finance firms, particularly in the renewable resource area that is otherwise receiving a lot of attention at the moment.

Rood: That is peculiar. The observation here in the financial industry is that we are seeing an overload of retirement funds that now are on the hunt for long-run projects with clear descriptions (understand what you are investing in) and that this may benefit open (fiber) networks (that reduce risks).

Just remark that in a climate of all kinds of financial trickery investors see money evaporate for so-called AAA investments.

In contrast - funding an open fiber network has

1. An immediate capital demand

2. Long-run, not short-run returns

3. When sticking to passive nets, you know precisely in what asset you invest

4. When open you have the advantage that it could not easily be duplicated and reaches high penetration.

5. The default risk of having active operators on it is easily taken by new firms that use the net. And as a side effect, it provides some construction jobs that benefit the lower economic strata of society, thus reducing mortgage defaults

This is different from a closed optical network going in direct competition with incumbents, that creates

1. An immediate capital demand

2. Long-run, not short-run returns

3. Due to vertical integration it goes into head on competition

4. Penetration level in head-on competition is risky. Reflected in huge premiums, not quite AAA, more BBB

5. Risk of default is for the entire optical and optoelectronic infrastructure. Bargain hunters could then re-enter on the cheap.

Consequently High Return High Risk is the Name of the Game

Now if you are not allowed to put money in high-risk, and all those AAA financial constructs evaporated from the market. Where do you put next months millions of incoming retirement payments?

COOK Report: You are probably in better condition in Europe than we are in the wild and woolly USA.

My understanding is the credit seizure is that our whole banking system suddenly has no idea how much its underlying assets are worth. The major reason is that with the "free market" we have seen the emergence of financial instruments that have mushroomed that slice, dice and repackage and resell the debt to the point where no one quite understands who has what and what their portfolios of SIVs and CDOs are worth. Then because they don't know what they are worth, they are also illiquid (can't be sold). The tech bubble followed by the housing bubble has put the global

banking system at risk because in part the assets are in things like hedge funds that are outside the traditional banking system.

I find John Mauldin's internet newsletter service fascinating and far more informative than anything in the mainstream media. Since Thanksgiving (November 23, 2007 - Mauldin has been extremely worried.)

<http://www.investorsinsight.com/otb.aspx>

Odlyzko: Yes, he may be right, but it may take a while because of the (almost certainly unrealistic) expectations that appear to be held by investors.

As just one example, suppose we take the cost of FTTH at \$2,000 per home (rather high, but let's be conservative), and have to pay 5% for the money (about the high rate corporate bond yield right now, substantially over the cost of tax-free municipal bonds). Then the annual interest rate comes to \$100, or less than \$10/month. (Of course, there will be depreciation, ..., to come on top, but let's keep this simple.) But if you are an equity investor and expect to get 20% on your money (unrealistic, yes, but the worldwide rate on invested capital today is around 15%), you are talking of over \$30/month, a much

higher hurdle.

It's all a big puzzle, and I hope it does not get resolved in a gigantic financial debacle.

St Arnaud: I have talked to many carriers and potential investors in regards to FTTH. The biggest issue is not the return on the investment - it is the uncertainty of getting any return at all. And as we know the higher risk, the higher expected return.

North America telcos don't want to invest in FTTH (Verizon being the exception) because of the risk that most of the potential revenue will be captured by what they call "over the top" providers like Google etc. This is why they want regulatory relief in order to charge companies like Google to use their networks.

Most FTTH business cases assume the network operator will capture all of the cable TV and voice revenue. If that disappears because of new applications and over the top providers, the business case gets pretty weak.

Odlyzko: Well, these are not separate things. If you have a potential investment that has a 75% chance of getting totally lost, and you want to get 20% expected rate of return, you will be looking for a

rate of return of 80% in those cases where the venture actually takes off (plus a bit more for risk). If the expected rate of return is 5%, you are looking at 20% for the business to pay in those cases, 1 out of 4, that it is viable.

My strong suspicion is that with an expected 5% rate of return, you could have a net-neutral network infrastructure, say dark fiber, or wavelengths, even if most of the return were captured by Google and other companies in that sphere. But it is indeed hard to make a case if you expect 20%, unless you can charge according to value of communication.

Savage: OK, I'm coming in late here, but what effect, if any, does building a municipal fiber infrastructure versus a privately funded one have on this analysis?

On the one hand, if the construction is funded by bonds or at least with municipal backing, the risk of default is tiny, so the banks get their money no matter what.

On the other hand, if the project becomes a disaster, then it is the taxpayers who get hosed.

On the other other hand, irrespective of how it's funded - private or muni - the lower

you can drive down the cost/home the better off you are. So an offhand assumption of \$2,000 will really hurt you if in the real world it's \$1,000 or \$4,000.

Just some random observations at this point...

Vint Cerf: This strikes me as an argument for government-sponsored infrastructure that does NOT seek high rates of return but seeks compensation to recover maintenance costs. The interstate highway system generally falls into a similar category unless I am over-simplifying.

Savage: I respond with a resounding "maybe." The money to build the infrastructure still comes from somewhere -- taxpayers in the government model. And if it turns out not to be used (think "bridge to nowhere") it's still wasted. It's just that the waste is, or can be, hidden.

Don't get me wrong. I'm an information connectivity junkie, paying Verizon for the fastest FiOS I can get. But I'm old enough and wise enough, I think, to realize that I'm not necessarily typical

Brodsky: Could one of the ways a carrier generate a better return on its FTTH investment by selling [access

to other ISPs] to the network? That's a revolutionary concept these days, but I think there's a case to be made that line sharing and wholesale access could help their business case, rather than hurt it. Back to the future.

COOK Report: Hah!

I agree with you Vint. But has not fealty to the Chicago School been such that this would impossible to do? I mean if the "gummit" meddles in the free market its a fate worse than death!@! ! (to the Chicago boys).

Business Case Complexity

Waclawsky: Increasing bandwidth, increasing connectivity (eventually all people connected nearly all the time) leads to increased experimentation that in turn leads to the discovery of things people want to do (and will pay for). We don't know what they are but the short history of the Internet shows this experimentation engine (called the Internet) has generated enormous wealth. All signs are this will continue ...except for perhaps IMS or SIP ;-) ...so how does a business model forecast the future? I don't think the total impact of the highway system was well understood at the time and now

look at all the business it would later start, influence, augment etc. In a sense it is the same kind of enablement that broadband connectivity seem to offer.

St Arnaud: Another point I forgot to mention was "tyranny of the takeup".

It may be possible to get 5% financing for FTTH if you have 100% takeup. But if you are trying to deploy FTTH in an existing neighborhood in competition with cable or DSL (or both) then you are unlikely to get more than 10-20% take up per year. And if you don't offer cable TV, your take up is even less.

Low take up rates dramatically alter the economics as the average cost per house for FTTH, in the first few years, jumps from \$2000 to \$6000. I have lots of additional information on my blog - and actual business plans on my web site if you are interested in the details.

<http://free-fiber-to-the-home.blogspot.com/>

So to my mind there are only 2 possible financing solutions for FTTH: (a) government financed as Vint suggests (b) novel new business models (but as yet unproven)

The former may ultimately be the best solution, but I suspect it will be years before we

get consensus to move in that direction

Mark Cooper: The over the top argument is bull***. The marketplace should set a rate for the transport of bits that covers the cost of the network including a normal rate of return. The Telcos and cable operators want supra normal profits because they are not used to living in a competitive world. Of course, if you assume there will be little or no competition and no regulation, they will set their rates to achieve supra normal returns. They are already doing so, as the *Economist* pointed out in a January 17, 2008, entitled "Open Up Those Highways" lays that cuts through the bullshit "A New Yorker who wants the same quality of services of broadband has to pay around \$150 more per month" and the French are well behind at least a half dozen other countries. Unless you believe that the France Telecom is substantially more efficient than Verizon, this demonstrates the exercise of market power.

The recording industry made the same mistake in the 1990s. CD's were selling for \$12 (before the companies colluded to drive the price to \$17), because the production and distribution of CD's cost about \$8. When digital production and distribution came along, we should have seen a

precipitous drop in the cost of music. The Record labels looked at the market and said, \$17 is the value of the CD (even though our costs are only \$4). They wanted and expected the supranormal profits, but along came file sharing. The Record labels could not figure this out themselves, so then came Apple with iTunes. They look on Apple as an "over the top" seller, stealing their value, little better than the peer-to-peer pirates who started the whole thing. All for labels have broken with Apple in an effort to regain control over pricing (although in order to offer a product that might be more attractive to consumers they have unwrapped their content from Digital Rights Management software).

The record labels sold over 800 million singles last year at \$2 per single. They want to sell 400 million more CD (assuming people buy 2 songs per CD, which is supported by empirical evidence). The Record labels wanted to collect \$6 billion more for CDs, but consumers spent only \$1.6 billion on singles (and iTunes took half of that). The over the top innovators will steal the eyeballs, but they will still have to move the traffic. The network operators are making the key mistake of trying to tax the value of applications and content, rather than just rationalize the transport

business.

Coluccio: (Editor: referring to Odlyzkos' comment these are not separate things several paragraphs above). Andrew stated: "My strong suspicion is that with an expected 5% rate of return, you could have a net-neutral network infrastructure, say dark fiber, or wavelengths.."

This argument is bolstered, in my opinion, by the absence of back office system costs, which tend to be disproportionately large when compared to the over-the-toppers. In a purely wireline model, sans content, content delivery systems, CRM, NMS, billing services -- in effect, mountains of horrendously costly OSS baggage -- become the burden of the "over the tops", who, I'd go so far as to suggest, tend to be much leaner in these areas from the outset. This cost avoidance, going in, allows more revenue to flow directly to the bottom line.

It's a much leaner form of existence, I grant you, but one that, if played by monopoly constructs designed around an optimal number of wireline providers per demographic begins to make sense. Compound the foregoing with profits from resale to CLECs and Specialized SPs, who collectively avail themselves usually as profitable

customers for incumbents' wholesaling operations, and ... did I just hear someone say back to the future?

Waclawsky: This goes back to a discussion I started on the list six months to a year or so ago. How much of the operator infrastructure supports the bearer plane (the dumb pipes part) and how much supports the control/billing/management/etc. plane? Some people I have spoken to in the industry have given me estimates of network cost anywhere from 5% to 30% for bearer infrastructure and of course the rest is for the control/billing/management/etc. infrastructure. It seems that all the effort to monetize all the potential uses of the asset (connectivity), is much more costly (and complex) than the asset itself.

Odlyzko: A very good question, it would be nice to get good numbers for modern networks.

And don't forget marketing costs. (Perhaps you are including those in control/billing/management/etc., but they should be treated separately.) Most companies spend in the range of \$200-500 per new residential customer, and with substantial

churn, that adds up.

[Editor: and responding to Bill St Arnaud on January 21] **Andrew Odlyzko** writes: Yes, absolutely, that famous "tyranny of the takeup" that has bedeviled all telecom planning for the last two decades. When added to the presence of an incumbent that can lower prices at will, and the expectations of investors that they are entitled to Google-like returns on their money, it makes it very hard to do anything.

Thanks a lot for the URL of your blog, it is a basic resource that should be cited more often.

Jaap van Till: The key difference in take up percentage is made by preliminary (before any rollout) 'community demand bundling'. In a spectacular example in NL (OnsNet in the city of Nuenen) the citizens convinced each other and their neighbors to join and sign in for a community owned FttH infrastructure. They did this to get a lower cost of connection themselves and they did: 90 percent takeup! We see the same thing happening in our FttF experiment here: massive positive response at meetings with the farmers. Value chain reversal at work :-))

Vincent Dekker: And that (community demand bun-

dling) has have proven to be a good model here in The Netherlands. Reggefiber, the company that does most of the ftth so far, simply says to people asking for ftth. If you hand over a list with the signatures for some 40 percent of the homes in your town or part of city, we'll come in and do ftth for you. It works! People start lobbying and manage to get 40 percent or more before the first fiber is in the ground. For example in Nieuwland, 5200 homes and part of Amersfoort (one of the ten biggest cities in The Netherlands), Reggefiber started deployment in November last year and had over 50 percent market share to start with.

See <http://www.nieuwlandopglasvezel.nl> Subscribers get 50 mbps up and down, cheap phone calls and a lot of TV for 19,95 euro (30 dollars) a month the first six month and after that for 50 euro a month. Even the 50 euro are probably less then what they pay now for much less speed and fewer TV channels and their phone bill. So who wouldn't want that?

Cole: Such "demand assembly" is one of the ways communities can help persuade infrastructure builders to deploy, if the builder is willing to consider such data. It was a

strategy that included this in part that the City of Fort Wayne, Indiana, USA used to persuade Verizon to jump Fort Wayne in its queue. Unfortunately, Verizon has now signaled that it is not accepting such guidance, at least until it finishes its current build-out plans (whether it could be persuaded to change its mind is another thing).

The other potential players in the US, as many on this list can testify better than me, face a variety of regulatory constraints imposed in part with the support of "dog-in-the-manager-like" statutes that if the incumbent is unwilling to deploy, it must be a "bad" deal that the city should not get sucked into.

In the staff report to the city of San Francisco, both the cable company and the telecom are quoted as saying that San Francisco has all it needs (or is willing to pay them to provide, which to them is the same thing), and therefore the city should NOT be considering fiber. **The staff report implies that "demand assembly" would NOT sway either incumbent, and unfortunately it is a costly process to do the demand assembly to test that implication.**

One of the other ways to "assemble demand" is to offer the local government as an "anchor tenant" which might be enough in itself to tip over the decision. Ironically, in the San Francisco case, the city has already "fibered" itself, thanks in part to use of incumbent conduits where the permission was conditioned on use by the city only.

I still hold out hopes that some clever "urban renewal" proposal might pick either a depressed area of a larger city or a smaller depressed town, and see if at least some of the dreams of the "technology visionaries" can be demonstrated -- rising property values, improved delivery of public and private services, etc. For instance, in an experiment in "bread and circuses" or to use the later term "opiate of the masses" would close to 100 percent delivery of a triple-play package to a crime-ridden public housing project reduce crime and otherwise improve the living conditions in the project? It strikes me that this experiment, complete with some minimum-quality computer/TV in each unit might cost far less than some of the social experiments already launched for such places.

Esme Vos: I heard Viviane Reding speak on a panel at DLD in Munich early this week and indeed she's very

keen not only on improving broadband penetration and speed across the EU, but also in harmonizing spectrum rules. Ms. Reding once again cited to the Netherlands as a broadband paradise like Sweden and Finland. Alas, if only it were true.

Here I am in the center of Amsterdam, and I cannot get more than 20 Mbps downstream / 2 Mbps upstream from any of the DSL operators. There are many of them but this appears to be a limit. When they say 20 Mbps, you know it's half to less than half of that.

I think it's time to STOP talking about the Netherlands as some kind of broadband paradise because I am still stuck on ADSL, thanks to KPN's lousy copper network, which is used by all the other operators. And I live in the center of Amsterdam, not some tiny village near Enschede. Bah!

So I am not going to listen to any more useless babble about the Netherlands being such a great broadband place. Get me the 50 Mbps+ symmetrical bandwidth, then I'll believe it. I do not understand why Reggefiber and all those supposedly amazing FTTH companies don't set up infrastructure here because here in Amsterdam is where the money is.

Dirk van der Woude: Esme sadly is quite correct. Hartenstraat is quite far away from any CO and next to that some of the copper under the non business parts of Amsterdam's canal district may well have neolithic origins.

By the way, Esme, same goes for Denmark: hailed as a BB paradise - until one learns that, over there, even less than the FCC 200 K standard is counted as BB.

I do hope to have some news for our citizens in the yet FtTH challenged boroughs soon. Only not yet, alas.

Conal Henry: There is a price that the consumer is prepared to pay for broadband - it's the price they have always paid for their telecommunications - here in Ireland its about €50/month for fixed line. The consumer will spend this on the best available offering but will not spend more to make a better offering available (my thesis - not fact).

The benefits of high speed, high quality communications infrastructure accrue to different stakeholders in different ways. For the average telco - There might be €50/month in it but that's not really driven by the quality merely the access. There is a hope that some spend from other markets (eg broadcast or even retail) can be

brought in - but this case is still very weak. For the consumers there are lots of benefits - but they are not readily understood pre-purchase.

Richer communications, a wider field of experience, richer access to the rest of the world, improved communication and information management - all things the consumer will enjoy when they have them but not things they tend to miss when they don't.

For the state there are a myriad of benefits - Better connected citizenry generally make more capable citizenry and therefore more likely to earn the state money and make the state more competitive globally.

Better communications infrastructure reduces reliance on and cost of other expensive infrastructure - such transport, power, education, broadcast Better communications infrastructure reduces the cost and improves the efficiency of many state activities - taking taxes, communicating with citizens, providing resources (welfare, health etc)

So we have a situation where the consumer wants it but doesn't really see it, the telco has to do more to get the same - but where the state can dramatically improve the

national well being.

The same has been true in the past of investments in roads, running water, education, electricity distribution, postal communications, And many other areas - more often than not (though not always) the state is the best investor in these areas.

I believe that those countries that get this and invest for the state benefit will stride many years ahead of those countries that try and build their infrastructure off telco and consumer benefit propositions alone.

In the 20th Century those countries with the most up to date industrial infrastructure grew fastest and strongest - in the 21st century it will those with the strongest communications infrastructure.

Craig Dobson: A year or so back I developed the economic models for a variety of bb deployment strategies for a vertically integrated provider and indeed, as will be no surprise to those here, the RoI numbers become increasingly negative as fibre levels increased. The issue was not only the actual deployment costs on a homes-passed basis with low take rates but also the costs to upgrade customer home infrastructure and provide STBs (both free

of charge) to those that did take the services. When the plan is to parcel out the bandwidth so it can be monetized, the biggest bang for your dollar comes from upgrading copper.

Subsequent discussions around ways in which fibre deployment could be increased by sharing the deployment costs amongst various players via novel ownership, unbundling, wholesale, etc. arrangements went nowhere. At least with the incumbents I've dealt with, infrastructure is a trump card and there is no appetite to open it up and chance not being able to compete on services.

The issue around customers not knowing what they don't have is significant and education is expensive. Both on the private and public side, arguing for fibre infrastructure that will support a myriad of to be developed services is a tough sell against the 'risk-free' free 'broadband' infrastructure offered by the incumbents, particularly when the incumbents won't necessarily agree to providing services over infrastructure they don't own.

It seems to me that compounding this is a chicken and egg problem in that **arguing for or establishing a structurally separated network in a community**

requires an established ecology of service provider businesses if the community is to realize the benefits of services-based competition. Given the ecology won't develop in the absence of the network and it's tough to justify the network when the ecology doesn't exist, how do you kick-start the process?

Wireless: Seeing Some Good Deployment in 3.65 GHz

Harold Feld: Sascha Meinrath (who is also on this list) has put together some "notes from the field" for early deployment in the 3.65 GHz band. Given that this band was only opened 6 months ago, this is very promising for deployment. You can see Sascha's posting: <http://saschameinrath.com/node/564>

What we will need to see, as deployment increases, is whether the FCC's mandatory voluntary coordination approach works. (The FCC requires folks to coordinate, but leaves it to the folks on the ground to work out the coordination themselves. That's why it's "mandatory voluntary.") We also need to see what happens when the relevant IEEE committee comes out with appropriate standards. I think most of the initial deployment is using WiMax contention based protocols, which limits them to the bottom 25 MHz range. Finally, it will be useful to see what level of penetration it takes to get mobile chips put into devices under the FCC's hybrid approach. Right now, I believe the use is all for the fixed "high power" base stations.

We'll know this band has really made it when one of the big boys like Cisco buys one of today's entrepreneurial companies and starts cranking out equipment by the boatload.

Nevertheless, the fact that we are seeing first generation equipment deployed after only 6 months is pretty impressive and a strong indicator for how fast this sort of technology can move, especially when compared to licensed networks. I make this observation in my own blog (natch): "How To Give America A Wireless Broadband Network for Christmas 2009." <http://www.wetmachine.com/item/1029>

COOK Report: Great post and Sascha's blog is well worth the read, and this is very relevant to 802.11y of which Peter Ecclesine is the chair and is also on list.

Two days later: Had a detailed conversation with Peter just now. He will join Sascha Meinrath on Monday evening the 24th at the "Cook-in" to talk about 802.11y and the 3650 Ghz band. 802.11y offers a control system for

wireless connection of your device to compatible infrastructure.

I took notes and sent to Peter asking him to clean them up and send back. What follows is a tad dry compared to my rather awe struck reaction as i listed to what he had to say. This is exciting stuff that will give us a whole new way of thinking about how wireless fits into the picture.

Any way here is Peter's summary ---

"As Wi-Fi radios evolve, they tune more widely into adjacent spectrum, like from 5 GHz to the 4.9 GHz Public Safety spectrum and the 5.9 GHz Intelligent Transportation System spectrum. What has been missing is the software/firmware for licensed operators to control the radios.

These radios running under the FCC lightly licensed scheme will use an over the air control system; listening for a local infrastructure signal that enables them to negotiate link and function responses. They don't operate

until they receive an enabling signal to begin operation.

This url for a late 2006 Cisco FCC meeting explains a good bit:

http://fjallfoss.fcc.gov/prod/e/cfs/retrieve.cgi?native_or_pdf=pdf&id_document=6518438394

802.11y is mostly a software protocol that can easily be loaded into 802.11 chipsets numbering in the hundreds of millions per year (11a/11n). For 3650 MHz operation, additional Energy Detect features are specified, to enhance operation in the presence of other radios.

The same 'enablement' protocol can be applied to other lightly licensed bands, with additional sensing and location-awareness requirements from the regulators (like the TV-band white space investigations).

IEEE 802.19 Coexistence Assurance has been studying operation in the 3650 MHz band, and their current modeling documents are found at

<http://ieee802.org/19/pub/2008/07/0011r13> Parameters for Simulation, and 07/0020r3 Coexistence Metrics for the 3650 MHz Band. (note that simulations will follow after there is agreement on the parameters and metrics)"

The Economics of Triple Play, Bundles, and Other Marketing Ploys

Mark Cooper: Verizon emphasizes triple play, so the high ARPU customers are the bundle customers. **Also, when a customer goes triple play, Verizon stops maintaining the copper, which, if they can get enough people to switch, allows them to lay off a lot of labor. A lot of cable only customers or high-speed only customers is a disaster for Verizon. The need a high (e.g.50) percent take rate for all three services to make a go (video and high speed on the revenue side voice on the cost side), which is why they have fought so hard against build out requirements.**

In Montgomery County Maryland, one of the highest average income counties in the U.S., they planned to roll FIOS out to only 60% of the country and only agreed to a broader deployment (down to 15 people per square mile) after a judge ordered them to arbitrate. This suggests that given their business model, over half of their service area is not economic.

Cole: At the Killer App conference in May 2007, I heard fairly major Verizon folk

make the same point about needing TV to make it all pay. (They were quieter about replacing copper, but did not say anything inconsistent with that.) In fact, they were stressing all kinds of new "services" such as music videos, various kinds of video on demand, etc. all designed to extract additional revenue per site.

That is why I find the financial analysis that suggests benefits to dividing into a loop-only component (among others) so interesting. The Verizon folks did not believe that any loop-only organization could work, even if it were the "unfair competitor" of something entirely publicly-owned.

The other development that is sooo fun to watch is the provision of current TV programming via abc.com, fox.com, etc. ABC.com is even experimenting with making HD content available that way (i.e. most of its "Lost" series episodes, plus sample episodes from others). They claim it will work on a 2M connection, but my 10M had some jitters.

Hulu.com and other aggregators are also beginning to make current and past TV available, both directly and as a portal to the network (pun intended - sorry) web sites. At some point, one can imagine an offering that is "data"

only with the idea that if you want TV, get it from the TV web sites.

A reminder of the Chris Savage question -- what happens when voice is carved out of telecom and TV is carved out of cableco. Do the firms die, and if so, who else is injured in their death throes?

Felten: This is exactly why telcos should spend more time on their core transport trade and make it work out. What they provide has value, undeniably, so they should not die, they should be able to monetise it for all players in the value chain, from the customer who gets a free for all to the ad financed 'free' service who uses their network.

I know this is a contentious subject, but the stake is really high, and the ecosystem is very unstable at this stage...

Goldstein: There's more than one way to monetize a network. The Wrong way is for the carrier to become the shipper, or demand a kiosk arrangement, in which the user of the network pays based on somebody's perception of the value of the content. The Right way is for the carrier to charge enough to make a profit, and let its users figure out how to make

that pay.

The model that ATT wants, based on intense DPI and IMS, would for instance charge a few cents per email, a percentage of each e-commerce transaction, a bit off the top of a banking transaction, a share of the price for downloading things with a minimum to pay for pesky freebies, etc. Sure, that would "monetize" the net... and kill it. That, and not a "two-tier Internet", is what is behind the NN debate.

Tim Cowen: This keeps bothering me so I went back to some basics. Information goods are goods that have differential value depending on the use to which they are put. Magazines, music, movies, stock data etc all have differential value depending on the time the use the audience etc.

Value based pricing leads to differential pricing depending on what different customers value and can pay. Versions of information goods are created to take advantage of this and make goods appeal to different segments which will pay different prices for the different versions.

Intellectual property rights (IPR) are used to ensure that value is trapped in different information goods, versions etc to maximise returns. IPR is neutral, how it is used is or

can be an issue depending on things like market structure and barriers to entry in the relevant product and geographic market. Remember the fuss about video killing the radio star and TV industry filing suits to prevent home copying of TV programmes? Hollywood now probably makes more from versions and DVD sales than first offer box office.

The 'wrong' and 'right' way referred to above is dependent on whether there is or is likely to be a competitive market; if there is, then prices will be eroded and competition will bring a more reasonable limit to the profit than would otherwise be the case. If no market, then monopoly rents will apply absent regulation, and innovation will suffer.

The issue with AT+T is that the monopoly area is not well defined, not well regulated and the risk of ability to act as gatekeeper or extract monopoly rents it is increased. But that is what you would expect to get with an ineffectively regulated monopolist.

Anyone disagree?

Goldstein: That's the point I'm making. In a competitive market, their plans would go

nowhere fast. As I wrote in mid-2005

(<http://www.ionary.com/ion-skyped.html>), the "fat wasteband" scenario of DPI'd "broadband" would never fly if there were competition at the IP layer, and **therefore the big ILECs have demanded (and gotten) forbearance from common carrier obligations. Thus there is no longer an independent ISP function. The wire owner controls the content. That's an ineffectively regulated monopolist.**

Intellectual property, however, is irrelevant to my argument. Yes, ATT is now conflating the two, and talking about putting "anti-piracy" filters into the bitstream. Truly dumb, and as Gordon notes, doing so removes some legal protections (not responsible for carried content put there by others) that Verizon, at least, understands. But Fat Wasteband isn't about protecting IP; it's about taking a piece of all of the action across the wire, including e-commerce.

Marks: A straw regarding triple-play - my Comcast \$100 triple play expired, and now we're paying about \$180/month. My wife's immediate reaction was "cancel the TV, the DVR hasn't recorded anything in a week, and do we need the phone along with the mobiles?" She

didn't consider dropping the net connection.

Peake: Last data I saw, I think November 07, there were 1.3 Million FiOS subs and 717,000 FiOS triple play. FiOS services are rolled out at different rates, ordinary FiOS often coming before the full range of TV etc. Churn is low, 1.3% month, and 70% of subscribers are new to Verizon broadband, so it's not Verizon DSL to FiOS migration. Agree with Benoit, FiOS seems to be doing well, and gets a good customer satisfaction rating. (Much as it hurts to be positive about an RBOC :-))

Cerf: Is it true that installation of FiOS to include voice causes permanent cutting of the copper twisted pair so you cannot go back to it for voice service or dsl?

Turner: Yes, it's true. In the case of my FiOS installation, I had analog phone from Verizon (and a cable modem). I signed up for FiOS to get higher speed Internet. The installer's standard procedure called for moving my voice line over to the new optical network termination box and then removing the (overhead) copper drop from the side of the house all the way back to the connection at the telephone pole.

I was able to talk him out of that by claiming (not strictly

true) that my niece in the back apartment still had an analog phone running over the six pair copper drop cable. But his procedure called for removing the former copper drop wires back to the pole.

Savage: You can always "go back" to copper for some cost. Aside from making the copper loop unavailable to competitors, I wonder if Verizon has a sort of "burn the ships so the army can't run home without victory" kind of idea here...

Cooper: Verizon chooses to cease maintaining the copper. The cutting of the copper is not a consequence of the pulling of the fiber, it is a policy choice on Verizon's part.

John St Julien: Mark Cooper noted that, based on Verizon's behavior in N. Virginia, the incumbent doesn't plan to serve more than 60% of even wealthy communities.

That's not only true of Verizon's relatively expensive FTTH infrastructure. **AT&T (back when it was SBC) was open about its policies told investors that it wasn't going to serve more than 52.5% of its installed base. (Or so I figured from how much of its "high," "medium," and "low value" customers it was planning to serve in light of the split between**

those categories in its installed base that it claimed in the same presentation.)

I'd say its pretty consistent over time: the telco incumbents in the US don't intend a full build-out. Or at least not any time soon.

Incidentally, though I didn't recall it when I went looking through my archives, Mark Cooper was cited in the same post where I worked this out. He's been sounding the alarm for a good while now.

Tim Cowen: I think we need to separate out monopoly and non monopoly businesses. There is nothing wrong with pricing in a way that shares risk and captures value in competitive markets for example in relation to complex communications transactions. Many companies do this in the complex communications and IT area and the idea that risk and reward should be aligned is uncontroversial.

Pricing monopoly assets according to what the market can bear is arguably illegal in a number of jurisdictions. **Access in communications has often been looked at as displaying TINA features: (Tina is a simple rule and stands for There Is No Alternative.)** Seems that it would apply to those

who live out in the sticks, or as discussed over the past few weeks, in locations where there is no competition. Limited availability of competition and bandwidth services are, and are likely to be, limited depending on geography. In a timing sense there may be competition even if not currently two or more suppliers to a given location if alternatives will be supplied on a timely likely and sufficient, if not readily available, basis.

In simple terms, if there are enough good choices of suppliers and products and services at competitive prices, then the market can be seen to operate. **Only where Tina applies will it be right to accept that there is a market failure and in that situation, regulation should apply. Regulation is then there to mimic the effects of a competitive market and drive investment (there is a very good Bobby Willig study of this in comms in the US). This is the basis on which we have accepted a suite of undertakings that regulate our Openreach access business here in the UK.** (However, in principle the facts have never been proven, and as a practical matter there is a long list of advantages that clear regulatory boundaries bring, so that it was also sensible to resolve the issue and move on). **The fact that Tina applies to**

access in the EU is not in question and is the basis for EU wide regulation and has been for many years.

In IT and telecoms it is important to remember that there are many types of bottleneck, and the most obvious being such things as those decided by the ECJ in the recent Microsoft and other judgments and Commission decisions. (I did a speech on this for the ABA a while back which I can dig out if anyone is interested.)

The issue we should really be discussing isn't whether companies can price and earn returns in competitive markets. That should be accepted. Also, we should not be discussing whether monopoly rents are against the public interest. We should agree both and while we may agree or disagree on the boundary the real issue is really why access regulation is so weak in some places and not others. The Openreach model is an application of access regulation with some reinforcement around incentive alignment and transparency. This is now becoming recognized as a way that access regulation can address the bottleneck issue and is happening in Ireland, Italy, Sweden, Australia and New Zealand. Nothing is perfect but the alternatives have been tried and failed. (The US is a

great example of the failure of structural separation).

The US with its notably lagging access speeds and relatively relaxed regulation, is similar to the position in Germany and a few other places where powerful players and regulation are more entangled. Effective regulation of monopoly is something we should agree on and unfortunately that seems to be less the case than it should be.

Collapse of linear tv and other things

Savage: What about the fact that traditional linear TV is going to be dead in the relatively near future? At some time in (I suspect) the not-too-distant future, those in the linear video transmission mode will realize that they can get more bang for the bit in devoting bandwidth to customized streams.

Sebastian Hassinger: Not sure if this was mentioned before I joined the conversation, but I believe that 2008 will be a watershed moment for the death of 'linear' tele

vision and the tradition TV network and production model. The writers' strike is now impacting *next* season's content, youtube.com's traffic has more than doubled since the strike started, and you have an army of creatives sitting around, surfing the net, bored and wondering how long their savings are going to last. A perfect storm.

Cecil: At some point the old models' methods of self preservation actually hasten their destruction. I don't think this is limited to TV.

Hassinger: That's the central lesson in Collapse, by Jared Diamond. Focused on societal collapse but in my opinion equally, if not more applicable to firms. <http://www.amazon.com/gp/product/0143036556?ie=UTF8&tag=lostjuly-20&linkCode=as2&camp=1789&creative=9325&creativeASIN=0143036556>

We've certainly seen a lot of the "do more of what once worked when it clearly isn't working anymore" from various industries over the past

couple decades. The recorded music, motion picture, television & telecom industries instantly spring to mind. Further back, any of the examples in *The Innovator's Dilemma* also tend to follow that pattern.

Savage: David Sloan Wilson, in his "Evolution for Everyone," has a wonderful metaphor for this: "dancing with ghosts." He's talking about a species highly evolved to a niche which then changes/disappears out from under it. The poor animals just keep doing all they know how to do - dancing like they used to, but with no one there. But it isn't adaptive anymore, and they go extinct.

Eight Generatives Better Than Free: Is there an IP Economy Lesson here?

Herman Wagter: A thought provoking blog post is found at <http://blog.futurestreetconsulting.com/?p=42>. It is called "Unevenly Distributed: Production Models for the 21st Century." A snippet is repeated below.

"It's not that YouTube is competing with you for dollars - it isn't, at least not yet - but rather, it is competing for attention. **Attention is the limiting factor for the audience; we are cashed up but time-poor. Yet, even as we've become so time-poor, the number of options for how we can spend that time entertaining ourselves has grown so grotesquely large as to be almost unfathomable.** This is the real lesson of YouTube, the one I want you to consider in your deliberations today. In just the past three years we have gone from an essential scarcity of video media - presented through limited and highly regulated distribution channels - to a hyper abundance of viewing options."

"This hyper abundance of choices, it was supposed until recently, would lead to a sort of "decision paralysis,"

whereby viewers would be so overwhelmed by the number of choices on offer that they would simply run back, terrified, to the highly regularized offerings of the old-school distribution channels. This has not happened; in fact, the opposite has occurred: the audience is fragmenting, breaking up into ever-smaller "micro-audiences". It is these micro-audiences that YouTube speaks directly to. The language of micro-audiences is YouTube's native tongue."

"In order to illustrate the transformation that has completely overtaken us, let's consider a hypothetical fifteen year-old boy, home after a day at school. He is multi-tasking: texting his friends, posting messages on Bebo, chatting away on IM, surfing the web, doing a bit of homework, and probably taking in some entertainment. That might be coming from a television, somewhere in the background, or it might be coming from the Web browser right in front of him. (Actually, it's probably both simultaneously.) This teenager has a limited suite of selections available on the telly - even with satellite or cable, there won't be more

than a few hundred choices on offer, and he's probably settled for something that, while not incredibly satisfying, is good enough to play in the background."

"Meanwhile, on his laptop, he's viewing a whole series of YouTube videos that he's received from his friends; they've found these videos in their own wanderings, and immediately forwarded them along, knowing that he'll enjoy them. He views them, and laughs, he forwards them along to other friends, who will laugh, and forward them along to other friends, and so on. **Sharing is an essential quality of all of the media this fifteen year-old has ever known. In his eyes, if it can't be shared, a piece of media loses most of its value. If it can't be forwarded along, it's broken.**"

For this fifteen year-old, the concept of a broadcast network no longer exists. Television programmes might be watched as they're broadcast over the airwaves, but more likely they're spooled off of a digital video recorder, or downloaded from the torrent and watched where and when

he chooses. **The broadcast network has been replaced by the social network of his friends, all of whom are constantly sharing the newest, coolest things with one another.**

The current hot item might be something that was created at great expense for a mass audience, but the relationship between a hot piece of media and its meaningfulness for a micro-audience is purely coincidental. All the marketing dollars in the world can foster some brand awareness, but no amount of money will inspire that fifteen year old to forward something along - because his social standing hangs in the balance. If he passes along something lame, he'll lose social standing with his peers. This factors into every decision he makes, from the brand of runners he wears, to the television series he chooses to watch. Because of the hyper abundance of media - something he takes as a given, not as an incredibly recent development - all of his media decisions are weighed against the values and tastes of his social network, rather than against a scarcity of choices.

Ed Pimentel: I suggest a look at http://www.kk.org/thetechnium/archives/2008/01/better_than_fre.php

Editor: Good recommendation from Ed. Here are a few snippets:

Better Than Free

The internet is a copy machine. At its most foundational level, it copies every action, every character, every thought we make while we ride upon it. In order to send a message from one corner of the internet to another, the protocols of communication demand that the whole message be copied along the way several times. IT companies make a lot of money selling equipment that facilitates this ceaseless copying. Every bit of data ever produced on any computer is copied somewhere. The digital economy is thus run on a river of copies. Unlike the mass-produced reproductions of the machine age, these copies are not just cheap, they are free. [SNIP]

Our wealth sits upon a very large device that copies promiscuously and constantly.

Yet the previous round of wealth in this economy was built on selling precious copies, so the free flow of free copies tends to undermine the established order. If reproductions of our best efforts are free, how can we keep going? To put it simply, how does one make money selling free copies?

I have an answer. The simplest way I can put it is thus:

When copies are super abundant, they become worthless. When copies are super abundant, stuff which can't be copied becomes scarce and valuable.

Eight Generatives Better Than Free

"Immediacy -- Sooner or later you can find a free copy of whatever you want, but getting a copy delivered to your inbox the moment it is released -- or even better, produced -- by its creators is a generative asset." [SNIP]

"Personalization -- A generic version of a concert recording may be free, but if you want a copy that has been tweaked to sound perfect in your particular living room -- as if it were preformed in your room -- you may be willing to pay a lot. The free copy of a book can be custom edited by the publishers to reflect your own previous reading background. A free movie you buy may be cut to reflect the rating you desire (no violence, dirty language okay)." [SNIP]

"Interpretation -- As the old joke goes: software, free. The manual, \$10,000. But it's no joke. A couple of high profile companies, like Red Hat, Apache, and others make

their living doing exactly that. They provide paid support for free software. The copy of code, being mere bits, is free -- and becomes valuable to you only through the support and guidance." [SNIP]

"Authenticity -- You might be able to grab a key software application for free, but even if you don't need a manual, you might like to be sure it is bug free, reliable, and warranted. You'll pay for authenticity. There are nearly an infinite number of variations of the Grateful Dead jams around; buying an authentic version from the band itself will ensure you get the one you wanted." [SNIP]

"Accessibility -- Ownership often sucks. You have to keep your things tidy, up-to-date, and in the case of digital material, backed up. And in this mobile world, you have to carry it along with you. Many people, me included, will be happy to have others tend our "possessions" by subscribing to them." [SNIP]

"Embodiment -- At its core the digital copy is without a body. You can take a free copy of a work and throw it on a screen. But perhaps you'd like to see it in hi-res on a huge screen? Maybe in 3D? PDFs are fine, but sometimes it is delicious to have the same words printed on bright white cottony paper,

bound in leather." [SNIP]

"Patronage -- It is my belief that audiences WANT to pay creators. Fans like to reward artists, musicians, authors and the like with the tokens of their appreciation, because it allows them to connect. But they will only pay if it is very easy to do, a reasonable amount, and they feel certain the money will directly benefit the creators." [SNIP]

"Findability -- Where as the previous generative qualities reside within creative digital works, findability is an asset that occurs at a higher level in the aggregate of many works. A zero price does not help direct attention to a work, and in fact may sometimes hinder it. But no matter what its price, a work has no value unless it is seen." [SNIP and finally]:

"Careful readers will note one conspicuous absence so far. I have said nothing about advertising. Ads are widely regarded as the solution, almost the ONLY solution, to the paradox of the free. Most of the suggested solutions I've seen for overcoming the free involve some measure of advertising. I think ads are only one of the paths that attention takes, and in the long-run, they will only be part of the new ways money is made selling the free."

Cecil: Brilliant. This man values abundance because he can imagine a new paradigm. He focuses on how change *creates* new value rather than fretting over how to preserve old value rendered irrelevant by the engines of innovation.

Berger: Yes, this is one of the most profound synopsis of all that is changing on the Internet and beyond. Kevin Kelly gives voice to the radical

COOK Report: Yes this is an extraordinarily well done piece.

Some of it has been said before by JP Rangaswami in his blog discussion of the Four Pillars of the Information Economy and in the kernel essay for Confused of Calcutta.

<http://confusedofcalcutta.com/the-kernel-for-this-blog/> This essay is 1358 words. It says what it would take most people 10,000 words to try to express. Read it if you have not.

As always though with these leading edge conclusions broader public understanding lags "waaay" behind. Does anyone have any ideas on what can be done about that problem?

Part of the barrier to broader understanding is that the

economy is far more than just information and those who do build and trade and develop physical goods for the time being at least have not much reason to trouble themselves about these issues.

While this analysis however is critical fare for everyone in the telecom internet space, the obliviousness of the phone companies to it remains extremely frustrating. One would think there is a huge opportunity lurking here.

Michele Bauwens: Thanks for referencing that great piece by Kevin Kelly. Scarcity and abundance have long been at the core of my own concern, especially about how to fit the essentially abundance logics of the immaterial domain, with the relative scarcity of material production.

One of the key questions I'm exploring is how to interface the non-reciprocal logic of peer production, with the world of material production, where it can not be fully applied, but can be 'informed' by it, and this is explored at length in this section, see <http://p2pfoundation.net/Catagory:Design>

Here are some interesting citations setting the stage for thinking through the issue:

open access to digital design - perhaps in the form a global repository of shared open source designs - introduces a unique contribution to human prosperity. This contribution is the possibility that data at one location in the world can be translated immediately to a product in any other location. This means anyone equipped with flexible fabrication capacity can be a producer of just about any manufactured object. The ramifications for localization of economies are profound, and leave the access to raw material feedstocks as the only natural constraint to human prosperity. - Marcin Jakubowski [snip]

Cole: Capitalism has figured out numerous ways to reward (and thus encourage) isolated individuals for extraordinary talent, effort, and/or luck through systems such as the patent process and the way fame works. It has, in my opinion, done much less (albeit not zero) to encourage (via reward or otherwise) the coordinated, combined efforts of hundreds or thousands of individuals to a common end.

So while a massively parallel trial-and-error system may work better for some purposes (computer operating systems perhaps), it has its own costs and risks and its own need for coordination. Whether the benefit to cost

ratio is better or not depends on both sides of the equation.

Also, much of the discussion assumes "polar examples" when the real world allows for (and sometimes exhibits) all sorts of combinations of isolated individual effort/talent/luck PLUS massively-parallel effort and talent and luck. (Think of the whole world that sprang up around spreadsheets, or around popular operating systems, whether "open" or "closed.")

Consequently I would like to see some more thought to how coordinate and pay the costs of the superior massively-parallel ("open") systems.

Cecil: On the other hand, one could look at the corporation as capitalism's answer to "encourage (via reward or otherwise) the coordinated, combined efforts of hundreds or thousands of individuals to a common end." (You might also call it "government", but let's avoid that digression for now). The point is, I think, that most corporations process serially. Most, therefore, are not paragons of efficiency. Thus, Dilbert.

What is beginning to evolve are newer forms of business organization that permit, allow, or blasphemously encourage parallel processing. This list, for example, could be held up as one such early

example. More refined forms, I'm certain, are evolving in the cyber-ether.

Cole: Erik -- the corporation is a nice point, albeit as you say, it usually does not work as "parallel" as it could. Also, I would add that it actually consolidates investment, rather than coordinating it.

However, it is possible to imagine a corporation set up to be parallel, or at least more parallel than is the norm these days. Indeed, industry-focussed horizontal integration can work this way -- think General Motors, which, in the past at least, has allowed its various car lines to compete with each other not only in nameplate and finish, but even in frame and motor. It has been a "modern" innovation to try to use the same frame and/or motor across multiple product lines.

Microsoft, for instance, has both Outlook and Outlook Express. It also has multiple ways to do the same thing in most of its office software, in part because it merges the contributions of separate teams. Levi Strauss has slightly different "cuts" of jeans for various parts of the US (Seattle is rumored to favor the broadest "beam" of any city).

Franchises are an interesting blend -- when they are abso-

lutely rigid in style and substance, they are only parallel in ownership/investment. But one could imagine the opposite arrangement, where the ownership/investment is rigid, but the style and substance of each outlet varies tremendously -- like a cross-industry conglomerate that was all the rage in the 1960's.

On the government side, federal programs often face this question of how much "parallelism" to allow -- state-by-state variation (50 laboratories of democracy), city-by-city variation, or sharp restrictions on variation, perhaps even limiting activity to full-time federal employees.

Nonprofits also face this spectrum -- one national organization, or a federation of state and local organizations, and how much variation to allow each chapter.

Cecil: Yes; I couldn't agree more. And your points are apropos. GM and others are adapting.

When it comes to corporate space, if yours is the model that contains all other models, basically you win. Contain can mean a lot of things - building a "market" for your products or ideas or abilities; it can also mean supplying or enabling. I think the greatest irony is that there is no containing; only supporting.

Relative to today's models, Google seems to get this, though others might argue they display impulses towards containment. They supply up, making use of "because" effects. Others will adapt. (thus MS + Yahoo vs. Goog's brilliant "offer of help", which, IMHO was pure genius PR move). Corporations, particularly telcos, will be the last to adapt; they just can't help it; the legal and regulatory incentives pulling the other way are simply too strong, and they seem to imagine they can force the rest of the planet to do things their way. At the same time, the fact that phone calls are no different than email - both are simply software applications + data - is simply too obvious. Something will give.

With the size of the struggle over control of the Internet - old v. new or insurgency v. counterinsurgency, there are times I get the sense we are approaching the corporate equivalent of the KT boundary. That doesn't mean I see doomsday; it does mean that I see evolution occasioned by chaotic and sudden movement, not linear progression.

Returning to your points about corporate legal structure, recall, corporations are, in the law's eyes "people". Right now there are a variety of corporate forms, each with

its advantages and disadvantages - e.g. S corp = taxed once; C corp means you can create markets for equity, but the price is double taxation. At the same time, there are new ways of raising money and new forms of banking taking place on line. Sooner or later the law will have to catch up as business models are struggling to be backwards compatible with the law.

What I sometimes wonder is whether the social networking platforms - and dozens seem to pop up on a daily basis - will morph over time (if they haven't already) - as forms of doing business. Much of what social networking attempts to replicate mimics behaviors, information flows, relationship management (both buyer and vendor), and affiliations we see playing out in the "real" business world. They might become, more or less, a "second life" for new organizations of individuals engaged in business. Several large corporations are already beginning to use these platforms for business purposes. That will only increase. At the same time, it may be that cell phones are destined to be our next currency - who needs paper anyway?

Hassinger: I think the most interesting aspect of 'parallelism' is that which allows

autonomous and semi-autonomous individuals and groups to coordinate loosely amongst themselves and act as a whole with much larger scale.

In essence, isn't this the model of the Internet itself, and largely the reason for its complete dominance over any other network we've ever known? Is it a surprise, then that this model would migrate to the other fields of human endeavor?

Maybe this decentralized, loosely coupled organizational model is the default 'starting point' in the environment of information abundance that KK is describing. Maybe that's why the monolithic, rigidly hierarchical organizations seem to be the ones experiencing the most distress in our era.

Berger: I just want to point out that the patent system isn't capitalism. Its monopoly enforcement by a government. It may be used by capitalism, but itself is really the opposite of a truly 'unfettered' free market. But it does show that in reality we need government (at least if it actually represents its citizens) to modulate capitalism.

In terms of corporations. I think that Erik is on to something in terms of it being a mechanism to harness individuals into collectives. But

it's all modeled on 19th century scarcity, command and control, etc. I've always thought that corporations were the first AIs. They were just made up of people, not computers. Their actions are greater than the sum of their parts (and sometimes lesser than the sum of their parts)

I consider the "Free and Open Source" efforts the first organizations that are post-scarcity based. Still very primitive but probably a lot to learn from how they work or don't work.

Bauwens: Thanks for the discussion on parallel processing, but these probabilistic production processes are just a small part, of the open design movement's logic. The key is separating design from production, and the design path follows the free software path, voluntary contributions combined with universal availability. What this does is at the same time dramatically cheapen the primary 'commodity' of production, while it also poses the problem of how to sustain the process.

An important insight so far, if you want open design, is that payment may be (and usually is) counterproductive, as it 'crowds out' the passionate motivations that make for voluntary contributions, and thereby also, the striving for absolute quality (as compared to for-profit's relative

quality aim). Therefore, it is to benefit-sharing that companies that depend on a open design companies look, a general support of the process of voluntary collaboration, rather than a conditional payment to individuals.

Of course, a large part of free software programmers do get paid, either pretty much in a 'unconditional' manner (free to pursue their production according to the logic of the project itself), others 'conditional'. The Oekonux group therefore distinguishes 'simply-free' software, only the license, from doubly-free software, the license and the voluntary process.

I have heard from an IBM employee the following: the company has saved 90% of its infrastructure cost through the adoption of FLOSS, from this sum, it pays 5% to its own employees (conditional production) and 5% to the free software projects itself (which would count as unconditional benefit sharing)

Why I am explaining this, because this changes the debate about finding payments for parallel processing, since it is by definition counter-productive in a context of commons-based peer production. Yochai Benkler also evacuates the issue of motivation: he says that where the threshold is low enough, any kind of motivation becomes productive, so that is no longer an issue.

The result is that peer projects are sustainable collectively, but not 'individually'. As long as the projects can mobilize new volunteers to replace those they loose, the project is sustainable. However, individually of course, it is not sustainable.

Yet society is getting so many benefits from it, innovation is becoming diffuse and social and to a large extent related to precisely these peer to peer dynamics of freely contributed collaborative work, that it is worth

thinking about a general basic income that would sustain it.

In the Middle Ages, this was done for about a quarter of the male population, the monks, whose 'spiritual productivity' was recognized as extremely valuable, and therefore subsidized by the whole of society. Similarly, as peer-produced social innovation is becoming more and more valuable, it will become and needs to become a process that needs to be sustained through a general investment, much like the telecommunications and transport infrastructure. I'm calling this "partner state policies which empower and enable the direct production of value".

Google and Enterprise Utility Computing Sneaky or Empowering?

Cecil: A pattern repeated? Incumbent technology the enemy of innovation? Could cloud computing render corporate IT surplus? Resistance to change? (From FierceCIO)

Google tries to sneak one past IT During the last year, we've seen Google slowly sneaking into the enterprise, with a series of solutions that bring the company's web-based consumer technologies into a corporate environment. Just this past week, we saw the latest such product from Google, an email security and discovery suite based on Postini technology. Despite Google's best efforts, however, we haven't exactly seen enterprises rushing to deploy Google's software, which has proven quite popular with consumers.

So, what's a Google to do? Why, sneak its software into the corporate world, of course. **The company's latest enterprise product, Google Apps Team Edition, is a version of Google Apps that's designed for corporate users but can be used without the IT department's knowledge.** "Current business versions of Google Apps are linked to an organization's Internet do-

main and therefore require IT approval and at least some degree of administration," Ars Technica writes. "Team Edition eschews this approach, and allows end-users to create sharing workgroups so long as the individuals in question have valid e-mail addresses within the employer's Internet domain." See Google tries to sneak "Team Edition" suite past IT help desk

<http://arstechnica.com/news.ars/post/20080207-google-releases-new-team-edition-application-suite.html>

Team Edition contains the standard core features of Google Apps, save for Gmail, as that service requires a degree of IT oversight and administration.

According to Google Apps senior product manager Rajen Sheth, "Google Apps has been, by definition, an IT project, and now we want to let people use it without IT involvement." Signing up for Google Apps Team Edition will allow registrants to see which of their coworkers has also signed up, which, in theory, promotes additional collaboration. Google emphasizes that this type of two-way visibility will allow work-

groups to begin collaborating with each other-apparently spontaneously.

There is, of course, a rather obvious fly in this particular ointment. **Sheth suggests that IT departments and administrators shouldn't be upset about discovering unplanned and unapproved implementations of Google Apps running on the corporate network because "[t]he IT department always has the option to sign up for the Standard Edition for free if they want to provide control over this. This is a solid, happy medium."**

Felten: I think that ultimately the question is irrelevant. Google has a good understanding of the overall market trends. They need to be a SaaS player if they want to be present on the enterprise space. They know that consumers are also employees and would like to import their apps into work. Therefore, they enable them to do so.

Cecil: Perhaps within the context of - Google is moving into new space and using an interesting work-around. **Viewed from a larger perspective, what I thought**

was interesting was walled garden (corporate IT) vs the "outside". That's not to say we don't need corporate IT - I couldn't get things done without them. On the other hand, if the migration is toward cloud computing where I can pretty much do it myself with a centrally aggregated entity who can bring greater economies of scope and sale to the enterprise, then it sort of mirrors things we've discussed at the physical layer - i.e. that technology empowers end users who pretty much want big pipes and the ability to interact with whomever without intermediaries telling them how to do it. Of course it's different and whatever is chosen, it will have to work, but in terms of parallel trends, it seems there's something here. Or maybe I'm just not seeing it. ;-)

Felten: I suspect the underlying reason is probably that overall, IT doesn't want to hear about Google, because Google = Privacy Concerns (and Google = Consumer brand, also).

So I'd say it's clever. I'm not sure the appeal of Google Apps itself is good enough, though...

Cecil: Not sure I'd agree. I carry a portable law library

and case files for about 40 different matters on my laptop (properly secured, though 99% of it is meaningless outside the context of any particular case). I used google desktop for the longest time to find things; it was great. Along comes IT and bang! its gone. In its place - MS search client that just plain stinks. In addition, I use both skype and google /IM/voice to keep in touch with multiple counsel nationwide. This is especially helpful when we're running a zillion cases at once; phone calls are for the really important stuff, but IM is for air traffic control - it helps me keep things moving (and it keeps costs down because I'm more efficient & don't have to pay per minute rates for "long distance telephone" to my outside counsel - which, by the way, makes me crazy, considering what we're typically litigating).

Felten: You misunderstood.

There are amazing Google products that serve me in my work environment on a daily basis.

But Google Apps is a specific package with substitutes itself to MS Office, or at least that's my understanding. And my experience of these is that so far I'd rather keep on using MS. While I'd kill to have Gmail like functionalities on my current ultra stinky Lotus Notes, I'd rather have

excel for my spreadsheets that whatever Google provides. Their PDF reader is also below par.

Cecil: Got it. Agreed completely regarding Google Apps right now as a substitute for MS office. But that's going to change. Regardless, however, there's a trend. Thus, MS' offer for Yahoo, and, interestingly, Yahoo's anticipated rejection.

COOK Report: I have started reading Nichols Carr's the Big Switch: Rewiring the World from Edison to Google

Carr so far seems to argue that the network as computer is inevitable. Carr fits well into the context of my just completed interview with 3Tera a utility computing company that has developed AppLogic a utility computing OS that makes Amazon's EC and S3 stuff seem a bit primitive.

I am wondering whether there is not some reasonable balance to be achieved between the corporate centric expensive apps and the "I-plug-into-the-great-computer-in-the-sky" view of the world?

Anyone else reading *Big Switch*? What 3Tera describes by the way is QUITE powerful and so far as i can tell not well known,

<http://download2.3tera.net/demo/applogic20demo.html>

should give an idea to anyone who is curious. BT, by the way, is a 3Tera customer.

Robert Berger: The 3Tera stuff is pretty expensive. But it does presage a major trend. Its a good thing that Amazon EC2 and S3 are primitive. That is a feature not a bug. They are more like the IP layer of cloud computing. Unfortunately its proprietary right now (as is 3Tera), but its a good level of abstraction to build lots of stuff on top of it. Including stuff that make it easier to do higher order things.

There are other companies doing interesting layered stuff on top of Amazon (and at least one is designed to make their high level abstractions work on many different cloud infrastructures).

CloudScale where I am an advisor. They are still in early stage development (and looking for investors if you know of anyone interested :-)

And a newly announced Heroku (<http://heroku.com/>)

Hassinger: Well of course nothing is quite as cut and dried as Nicholas "IT doesn't matter" Carr would like it to be. But it is undeniable (and not particularly original to observe) that the network is increasingly the computer.

It's kind of inevitable isn't it?

I think the analogy to electricity and the evolution from having your own generator (and full-time engineering staff on salary) in the basement to plugging into the grid whenever you want it is an especially apt one.

It's also true that 2008 is a watershed year - virtualization of infrastructure is the number one investment area for Enterprise IT - whether using VMware, Xen, z/VM on big mainframes, or the new hooks for virtualization built into x86 cores and linux kernels, everyone's doing it, even Dilbert.

<http://www.unitedmedia.com/comics/dilbert/archive/images/dilbert20183362080212.gif>

Virtualization in the enterprise shops means a breaking down of the barriers between "my datacenter" and "some other datacenter" and "owning my IT infrastructure" and "getting my workloads done however, wherever." In other words, it goes a long way to breaking down the psychological barriers that have been artificially keeping the enterprise from capitalizing on the economics of cloud computing.

One critical unanswered question - what is the computing equivalent of kilowatt hours? That is, what precisely

is delivered when you plug into the wall? Or to put it more into telecom's terms - what is IT's dialtone? :-)

Today the spectrum runs from the very crude - Sun's CPU hours to the very sophisticated - 3tera's full-on app infrastructure, with Amazon Web Services' virtual servers and simple data tables firmly in the middle. So far my money's on AWS - it feels like the lowest-common denominator sweetspot that HTTP/HTML hit on in the first place, but I think this is still a wide open question.

Pimentel: Think about the ramifications of a successful implementation of Joel Adams MicroWulf using VMWare. <http://www.uberreview.com/2007/09/what-to-do-with-a-microwulf-joel-adams-will-tell-you.htm>

Or rethinking grids as in the case of Indiana University (naradabroker), which added Web2.0/SOA to their to remove all the complexity of Grids application development and infrastructure management.

http://www.thaigrid.or.th/wttc2007/files/20070604-02-mark_baker-middleware.pdf

Executive Summary

Utility Computing, pp: 1-5

This issue examines how increases in processing power and decreases in the cost of storage combined with affordable bandwidth on a global scale are enabling what is variously referred to as grid, utility and cloud computing.

While users still have their edge-based infrastructure and the client server model of the 1990s has, by no means disappeared, **the maturation of grid computing into utility computing offers companies of all sizes the opportunity to use the internet as a more powerful and less expensive alternative** to client server computing. In many ways with these events, the network is becoming the computer. User's machines are plugging, like intelligent terminals, into services offered by Amazon and Google and increasingly others that run on top of utility computing infrastructure and offer entrepreneurs ways to build services in software without having to buy and operate complex expensive hardware.

These changes underway at the level of the internet's "engine rooms" have profound implications. Are we, for example, looking at the emergence of a global computing utility that is becoming as important as the electricity grid? Also does the emergence of utility computing offer the carriers an inducement to transform their networks into data networks the primary purpose of which is to offer enterprises access to computing services of all types that could be safely and securely performed within the carrier's cloud rather than in house by IT staff at each separate enterprise. In such a scenario voice is no longer the major carrier product but rather a feature of all the data applications that utility computing enables. We shall see that there is some evidence that this the business strategy behind BT's 21CN, Global, IP Network's using 3Tera's ApLogic Utility Computing OS to run its own applications and then to offer services to BT customers.

(You have just read the introductory paragraphs of page one. the are a good summary of the introduction as a whole.) The reminder of the introduction discusses

Nicholas Carr's *The Big Switch: Rewiring the World from Edison to Google*, Carr shows how Google style computing is becoming an "open source" alternative to costly and proprietary enterprise data centers. Carr explains how all this is made possible by virtualization and then goes on to discuss 3Tera as the leading example of virtualization as an enabler of utility computing. The introduction then describes how BT has embraced 3Tera for itself and for its enterprise customers. It closes by raising the question of whether carriers can transform themselves by turning themselves from purveyors of voice into distributors of utility computing?

3Tera interview, pp. 6-17

An interview with Barry Lynnn, the CEO and Bert Armijo VP of Sales and Product Management describes the rise of grid, cloud, and utility computing differentiating between each. We then discuss the origins of 3Tera – a company that resulted from the merger of two independent efforts in the grid computing arena. As they explain what they were doing it was **start-**

ing out to solve the problem of how to scale online applications.

This led them in a direction where they developed grid tools that enabled applications to be ported to the grid without rewriting code. They called what they came up with AppLogic – an operating system for utility computing.

According to Armijo: AppLogic does not require any unique of specialized infrastructure. It is simply Intel or AMD servers, direct attached storage and gigabit Ethernet. There is no unique hardware whatsoever involved in building an AppLogic grid. We are trying to enable operators – people who already own large numbers of servers to add utility computing services to their product mix.”

AppLogic is available at 17 data centers around the world. It can be used in re-branded form by paying hourly fees to the hosting center. Then they have have small and medium business customers who sign up for the service directly with us. We direct them to hosting partners but we provide first line of support for them.

Most recently, we started offering AppLogic for the enterprise. This is our software with an enterprise-licensing model attached to it. In this

case instead of installing it on a metered basis we install it on a server-licensed basis.

Symposium Discussion In EU Regulation as Guarantor of Open Access p. 18

Editor: A society must be able to allocate finite resources for the benefit of its citizens. As I begin to grow more familiar with the long wave thought of Carlota Perez [much more next month], this strikes me as symptomatic of what she calls finance capital. We can't agree on a course of action because we are still under the grip of what she calls the installation wave. Although we are in the transition between installation and deployment, what is possible is still being constricted by the demands of finance capital.

To the extent I understand the Perezian point of view we must shift to a production capital mindset and the most likely way of accomplishing that is for the government to step in and take a larger role again. But what she calls intelligent government working in a core shared vision with business and society. I confess fascination as I am listening to her San Francisco talk Feb 21, 2008 and earlier today to her Amsterdam talk from November 2007.

EU-commissioner Viviane Reding (speech of January 16) **“The regulation of access networks keeps open the potential bottleneck to competition in broadband markets. European rules have provided a real choice for consumers.”**

“The European model is empirically proven to promote not just choice, competition and innovation but also investment. Redding goes on to say that European stimulus of investment has been very successful

Jim Kayne: **It's not as simple as "build it and they will come". Rather, it's the ability to generate flow. Impeded flow (read scarcity) stifles growth. You can't tell what will develop; whether it will be replicative or novel; primary, secondary, or tertiary; interactive locally, globally or in-between.**

On the one hand, if the construction is funded by bonds or at least with municipal backing, the risk of default is tiny, so the banks get their money no matter what.

On the other hand, if the project becomes a disaster, then it is the taxpayers who get hosed.

On the other other hand, ir-respective of how it's funded

- private or muni - the lower you can drive down the cost/home the better off you are. So an offhand assumption of \$2,000 will really hurt you if in the real world it's \$1,000 or \$4,000.

Just some random observations at this point... snip

Vint Cerf: This strikes me as an argument for government-sponsored infrastructure that does NOT seek high rates of return but seeks compensation to recover maintenance costs. The interstate highway system generally falls into a similar category unless I am over-simplifying. snip

COOK Report: Hah!

I agree with you Vint. But has not fealty to the Chicago School been such that this would impossible to do? I mean if the "gummit" meddles in the free market its a fate worse than death!@! ! (to the Chicago boys). snip

Waclawsky: **Some people I have spoken to in the industry have given me estimates of network cost anywhere from 5% to 30% for bearer infrastructure and of course the rest is for the control/billing/management/etc. infrastructure. It seems that all the effort to monetize all the potential uses of**

the asset (connectivity), is much more costly (and complex) than the asset itself.

Conal Henry: Richer communications, a wider field of experience, richer access to the rest of the world, improved communication and information management - all things the consumer will enjoy when they have them but not things they tend to miss when they don't.

Wireless p. 29

Very significant wireless developments on the horizon. Harold Feld: Sascha Meinrath (who is also on this list) has put together some "notes from the field" for early deployment in the 3.65 GHz band. Given that this band was only opened 6 months ago, this is very promising for deployment. You can see Sascha's posting: <http://saschameinrath.com/node/564>

Peter Ecclesine on 802.11y: **"As Wi-Fi radios evolve, they tune more widely into adjacent spectrum, like from 5 GHz to the 4.9 GHz Public Safety spectrum and the 5.9 GHz Intelligent Transportation System spectrum. What has been missing is the software/firmware for licensed operators to control the radios.**

These radios running under the FCC lightly licensed scheme will use an over the air control system; listening for a local infrastructure signal that enables them to negotiate link and function responses. They don't operate until they receive an enabling signal to begin operation.

Peter and Sascha will be at the cook-in on march 24 discussing a path involving the 3650 space, 802.11y and the emergence of optical photonics that may cause the difference between wireless and fiber to almost disappear.

Editor: More problem with the finance capital installation era

Tim Cowen (lead counsel BT): The issue with AT+T is that the monopoly area is not well defined, not well regulated and the risk of ability to act as gatekeeper or extract monopoly rents it is increased. But that is what you would expect to get with an ineffectively regulated monopolist.

Anyone disagree?

Goldstein: That's the point I'm making. In a competitive market, their plans would go nowhere fast. As I wrote in mid-2005 (<http://www.ionary.com/ion-skyped.html>), the "fat waste-

band" scenario of DPI'd "broadband" would never fly if there were competition at the IP layer, and **therefore the big ILECs have demanded (and gotten) forbearance from common carrier obligations. Thus there is no longer an independent ISP function. The wire owner controls the content. That's an ineffectively regulated monopolist.**

An IP Economy Lesson from Kevin Kelly p. 35

Ed Pimentel: I suggest a look at http://www.kk.org/thetechnium/archives/2008/01/better_than_fre.php

Editor: Good recommendation from Ed. Here are a few snippets:

Better Than Free

The internet is a copy machine. At its most foundational level, it copies every action, every character, every thought we make while we ride upon it. In order to send a message from one corner of the internet to another, the protocols of communication demand that the whole message be copied along the way several times. IT companies make a lot of money selling equipment that facilitates this ceaseless copying. Every bit

of data ever produced on any computer is copied somewhere. The digital economy is thus run on a river of copies. Unlike the mass-produced reproductions of the machine age, these copies are not just cheap, they are free. [SNIP]

Bauwens: The result is that peer projects are sustainable collectively, but not 'individually'. As long as the projects can mobilize new volunteers to replace those they lose, the project is sustainable. However, individually of course, it is not sustainable.

Yet society is getting so many benefits from it, innovation is becoming diffuse and social and to a large extent related to precisely these peer to peer dynamics of freely contributed collaborative work, that it is worth thinking about a general basic income that would sustain it.

In the Middle Ages, this was done for about a quarter of the male population, the monks, whose 'spiritual productivity' was recognized as extremely valuable, and therefore subsidized by the whole of society. Similarly, as peer-produced social innovation is becoming more and more valuable, it will become a process that needs to be sustained through a

general investment, much like the telecommunications and transport infrastructure. I'm calling this "partner state policies which empower and enable the direct production of value".

Google and Enterprise Utility Computing p. 42

The company's latest enterprise product, Google Apps Team Edition, is a version of Google Apps that's designed for corporate users but can be used without the IT department's knowledge.

Cecil: Viewed from a larger perspective, what I thought was interesting was walled garden (corporate IT) vs the "outside". That's not to say we don't need corporate IT - I couldn't get things done without them. On the other hand, if the migration is toward cloud computing where I can pretty much do it myself with a centrally aggregated entity who can bring greater economies of scope and sale to the enterprise, then it sort of mirrors things we've discussed at the physical layer - i.e. that technology empowers end users who pretty much want big pipes and the ability to interact with whomever without intermediaries telling them how to do it.

CONTENTS

The Rise of Utility Computing in the Context of the Global Grid

Could it Ever Offer Carriers an Opportunity to Cut the Umbilical to Voice?

First Some Background	p. 1
Connecting the Dots or Why JP Rangaswami is Smiling Tonight	p. 3
Utility Based IT Services to Global Enterprises – the New Carrier Business Model?	p. 4
Meanwhile Back at Home	p. 5

The Network Becomes the Computer: Grid, Cloud, and Utility Computing

An Interview with 3Tera Traces the Origins Functions and Benefits of an OS for Utility Computing

Grid, Cloud and Utility Computing	p. 6
Cloud Computing?	p. 8
Difference Between Amazon EC and 3Tera	p. 9
Genesis of 3Tera	p. 10
From Scaling Online Applications to a Grid Operating System	p. 11
The Infrastructure	p. 13
The Business Model	p. 14

Symposium Discussion January 16 - February 14 2008

In the European Union - Regulation Will Persist in Order to Guarantee Open Access

p. 18

Does the Credit Crunch Exacerbate American and European Tendency Toward High Return Short Term Investment?

Monetary Resource Allocation in the Credit Crunch p. 21

Consequently high return high risk is the name of the game p. 21

Business Case Complexity p. 23

Wireless: Seeing Some Good Deployment in 3.65 GHz

p.29

The Economics of

Triple Play, Bundles, and other Marketing Ploys p. 30

Eight Generatives Better Than Free: Is there an IP Economy Lesson here?

Eight Generatives Better Than Free p. 36

Google and Enterprise Utility Computing Sneaky or Empowering?

p. 42

Executive Summary

p. 44

Symposium & Interview Contributors to this Issue

Affiliation given for purposes of identification - views expressed are those of the contributors alone

Burt Armijo VP Product Development 3Tera

Michel Blauwens, Chiang Mai, Thailand, Foundation for P2P Alternatives

Robert Berger, Principal & Strategist, Internet Bandwidth Development

Art Brodsky, Communications Director, Public Knowledge

Erik Cecil, Telecom Regulatory Attorney

Vint Cerf, Internet Evangelist, Google

Roland Cole, Director of Technology Policy at Sagamore Institute for Policy Research

Frank Coluccio, President, DTI Consulting Inc. South Street Seaport, NY NY

Mark Cooper, Director of Research of the Consumer Federation of America

Tim Cowen, General Counsel for BT Global Services

Robert Cromwell, Attorney City of Seattle

Vincent Dekker, telcom report , Trouw, The Netherlands

Craig Dobson, Taylor Warwick Consulting, Alberta Canada

Harold Feld, VP media Access, Washington DC

Benoit Felten, Analyst yankee group and Fiber Revolution Blog

Fred Goldstein, Principal of Ionary Consulting, author of The Great Telecom Meltdown

Sebastian Hassinger, Stratgeist, Oracle

Conal Henry, CEO E-Net an Irish open access regional broadband infrastructure provider

Lars Hedberg, Director Swedish association of Fiber Communities

Barry Lynn, CEO 3Tera (utility computing)

Kevin Marks, now with Google, Coauthor Micro Formats, Quick Time Developer

Andrew Odlyzko, Director Digital Technology Center, University of Minnesota

Adam Peake, Senior Analyst Glocom, Japan

Ed Pimentel, CTO AgileCo, Alpharetta Georgia

Hendrik Rood is a senior consultant at Stratix Consulting, The Netherlands

Bill St Arnaud, Director Ca*Net4, Canada's high speed research network

John St Julien, fiber advocate and activist, Lafayette, Louisiana

Chris Savage, Partner Davis Wright Tremaine, LLP Washington, DC

Brough Turner, Co Founder NMS Communications

Jaap van Till, principal consultant and partner Stratix Consulting, Netherlands

Esem Vos, founder of Muniwireless.com,

Dirk van der Woude, Civil Servant Amsterdam and fiber expert

John Waclawsky, Chief Software Architect, Motorola

Herman Wagter, Director Citynet, Amsterdam

A Note from the Editor on the April Format and Presentation

This issue has an introductory essay an interview and 4 weeks of symposium discussion.

Text, URLs and Executive Summary: I have attempted to identify especially noteworthy text by means of boldface and red 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..

Coming in May - out about March 17 - Carlota Perez on the the place of ICT in completing the transition to a “golden age” of deployment and the replacement of finance capital by production capital. Excerpts from talks in Amsterdam (11 - 2007) and San Francisco (2 - 2008). Likely interviews for the summer months: Peter Ecclesine on 802.11y; British Telecom’s 21CN strategy, Taiwan’s National Digital Archive program. A June issue will likely not be out until about May 1.

Travel Schedule:

Feb 18-22 CUD meeting in San Francisco

March 18-27 VON in San jose and COOK-in in Loma Linda

April 5-19 Grid Computing and research on NDAP in Taipei

**On March 24-26 2008 in Loma Linda the first Infrastructure Cook-in.
You will find the Wiki and a list of confirmed attendees at**

http://www.socialtext.net/cookreport/index.cgi?cook_in

Thanks to BT, Cisco and Google for COOK - IN Sponsorships