



The COOK Report on Internet



© Cook Network Consultants

Ca*net 4 Plans Customer Control of Lambdas Regional Nets, Universities and Researchers to Be Able to Establish Wavelength Peering via Specially Tailored Switches Making Telecom a Customer Owned Asset May Create Favorable Impact on What Appears To Be “Surplus” Fiber Infrastructure New Tools May Enable Users to Solve Trust Problems & Build Communities

Many Emerging Internetworks Will Render Central Control Unlikely [Summary](#)

Our world has been clobbered this year. First by the economic collapse of much of the industry and then by the uncertainty raised by the September attacks. Starting with our “What Good is the Stupid Network?” essay, September 2001, and continuing last month with “End-to-end Principles and Trust,” we’ve been expending considerable effort to try to understand what these events mean for both the big picture and the long term. We are crossing many boundaries in our analysis in the hope that doing so can achieve some ‘outside-of-the-box’ insights. On September 23 we sat next to Natish, a young Cisco engineer, on a five hour flight from Singapore to New Delhi. We explained what we were doing in the draft of the end-to-end and trust article. Natish responded with much enthusiasm saying that he had never realized there were all these other implications to the work he was doing on ATM switches. We take his reaction as a sign that our exploration is useful.

The Internet and its transport technology will not suddenly disappear. The global economy cannot function without it. However, for the foreseeable future, the Internet will be subject to great stress as

those who see their interests threatened increase their efforts to control it. Such attempts to control are mostly harmful. This issue of the *COOK Report* will explore some of the ways that control layers are being built throughout the Internet and how emerging Internetworks built around new centers of trust can render such control ineffective. Last month, while we were expressing our concern for the continued viability of the end-to-end principles, Larry Lessig’s new book: *The Future of Ideas: the Fate of the Commons in an Interconnected World* was coming of the press. Lessig has come to many of the same conclusions that we have. In doing so, he bases the structure of his new work on network layers and the end-to-end arguments.

He offers persuasive arguments that those who would exploit the Internet will not be able to fathom what wins and what loses in technology if they do not explore its consequences in the broadest possible economic *and* social context. His book is eloquent if depressing testimony to the ability of threatened interests to divide and conquer their opponents. Lessig’s take is as bleak as our own was back in late June when we completed our “What Good is the Stupid Network” essay. A quote from that essay is the final quote in his new book.

In this issue, this exploration continues

Volume X, No. 10, January 2002
ISSN 1071 - 6327

with a review of *The Future of Ideas* and a discussion of two technologies that migrate control to the edge of the network and that in doing so give us a glimmer of hope in Lessig’s otherwise demoralized world. At the very moment when Lessig portrays how the judges have joined the politicians in building layers of control throughout the network and in roping off the commons in an onrush of greed, we also have the end-to-end principles under attack in all the ways enumerated by Clark and Blumenthal. It would seem that the barbarians are at the gates. Where then do we go from here?

On the Inside

Introduction	pp. 1-2
CA*net 4	pp. 2-18
Gerck	pp. 19-24
Lessig	pp. 25-31
ICANN	pp. 31-32
Highlights	pp. 33-38
Executive Summary	pp. 38-41

Some Answers

The last issue began to explore how the growth and development of the Internet over the last decade had increased the importance of trust. Let's explore these ideas further by asking, what is trust? It is one of those terms that everyone thinks at first they understand. We tend to see it in black and white terms. In reality, it is far more complex.

The article on the end-to-end principle, in last month's *COOK Report*, examined, with the aid of Einar Stefferud, a series of ideas involving the role of trust in the technical architecture of the Internet. Stefferud's comments were based on a model of trust developed in 1997 by Ed Gerck. We have now asked Ed Gerck who is CEO of Network Manifold Associates, Inc. <<http://nma.com>> to describe this model of trust for us. Gerck sees trust in the context of the engineering problem of communications.

In his words, "NMA's goal with these trust tools is to empower individuals and companies in controlling their own network connections. How they will use this empowerment to build their business and other relationships in their interactions is outside of our scope. How they will balance control at the edges with control from a center is also open-ended. The tools that we provide will permit them to make a full range of choices. Trust depends, among other things, on being able to choose."

Gerck's exposition goes beyond understanding trust in communication networks. It shows how trust can be induced (communicated) from human to machine, machine to machine, and machine to human. For trust is something essentially communicable. But there are rules that govern such communication. We need to identify and understand these rules before we can put them to practical use. Gerck's article explains these rules and shows how trust can be applied to a wide range of communication systems. The article begins on page 19.

At the same time in Canada Bill St Arnaud and CANARIE are ready to embark on a project that will ultimately put con-

trol over physical network bandwidth into the hands of the end users. CaNet*4 will result in the development of switches that will enable user control of light waves. They will use a variant of BGPto enable connection of light waves (peering) across the boundaries of autonomous systems, that is across separate networks.

At the recent annual Next Generation Network meeting, *Light Reading* interviewed Andrew Oldlyzko, who is now with the University of Minnesota's Digital Technology Center. Andrew said "Basic bandwidth is a commodity service." "You can fight it, but you're not going to win." According to *Light Reading*, Oldlyzko believes the money lies in providing "edge services" that hook consumers up to new applications using that bandwidth. See http://www.lightreading.com/document.asp?site=lightreading&doc_id=9529 Also Scott Clavena in a November 19th *Light reading* column titled *On the Crest of a Wave* described how carriers and even ILECs are beginning to buy lightwaves rather than light their own fiber. http://www.lightreading.com/document.asp?doc_id=9771 CANARIE is building tools to enable such edge services. Ca*net4 will test Oldlyzko's and Clavena's premises with a vengeance

It will take another year before we begin to have a clear idea of what the results will be with both of these projects (Ca*net 4 and NMA). Nevertheless, we find it fascinating that rather than being centralized as Lessig laments, power is being decentralized and placed in the hands of individuals. Lessig's new book shows why this movement to the edge and user control is the only alternative left to central control and single points of failure. While we are confident that both the trust and light wave technologies will work, what cannot be predicted is how they will spread.

We will begin with a discussion of the CANARIE CaNet4 developments by means of long interviews with Bill St Arnaud and Wade Hong, who at Carleton University has done the OBGp development. CaNet*4 is a high end research and education network with the initial

implementation of customer owned and controlled light waves very well restricted to a walled and non commercial garden. If the technology is developed and implemented as successfully as we think it will be, we find it hard to believe that commercial implementations won't follow. If they do, they will have profound effects on global telecommunications.

CA*net 4

Editor's Note: We interviewed Bill St. Arnaud and Wade Hong on October 22 and Bill again on November 2. In what follows, we have combined text from both interviews. Bill St. Arnaud is the Director of Network Projects at CANARIE. Since 1998 he has been responsible for the coordination and deployment of the world's first national optical research and development network, CA*Net3. We last interviewed him on OBGp in the November 2000 *COOK Report*. We also interviewed him more generally on CANARIE in the March 2001 *COOK Report*.

Wade Hong is an engineer (BScEE) with the Physics Department at Carleton University since 1991. He is Manager of HEPnet Canada which supports computer wide area networking in subatomic physics research across Canada. He is actively involved in various advanced research networking and scientific computing projects.

Project Implementation and Goals

COOK Report: When you issued your CA*net 4 Design Document back on August 21 you stated some really ambitious objectives. "The overall objective of the CA*net 4 network is to build the world's first customer empowered optical IP network where end users, initially at GigaPOPs, but eventually individual institutions and ultimately users at desktops will be able to purchase, setup and manage their own optical wavelength networks. Using a set of concepts called object oriented networking (OON) in combination with advanced high performance distributed applications called "grids" or "eScience," this architecture

will enable massive bandwidth interconnectivity to support applications in basic research and education that will be carried out not only between universities and research centers but also between schools and eventually within our communities.”

How do things look 90 days later? Is the vision and proposed implementation substantially still the same? How about the schedule?

St. Arnaud: The deployment of the network will very much be dependent on government funding. The funding will also directly determine the scope of the project and the schedule. We feel confident that the Canadian government understands the importance of CA*net 4the project. [Editor: Sources say that the hoped for announcement is likely on December 10, 2001.]

COOK Report: Given what we read about Industry Minister Brian Tobin’s views regarding the importance of an “innovation economy” for Canada, http://www.timecanada.com/story.adp?storyid=19&area=_toc I would think there is strong reason to be hopeful.

St. Arnaud: Yes. We are very fortunate in Canada to have strong political leadership that understands the importance and value of research networks to the future of our economy.

COOK Report: Meanwhile, your Background Information slides say that, on receipt of funding, you will purchase wavelengths, equipment for your GigaPOPs, and switches. From the conversations we have had, OBGPIs clearly the key technology underpinning for everything else that you want to do. Given that you still have a fair amount of development work to do on the switches to implement OBGPI, do you still intend to set up all the other infrastructure on receipt of funding?

St. Arnaud: Again, this largely depends on the amount of funding and time frame in which it is made available.

Initially we may have user owned wavelengths with no OBGPI. As you suggest, it

could take several months for us to work with vendors to implement, test and deploy OBGPI on their products. It is important to note that OBGPI is only one possible tool in a much larger toolset for achieving this vision. Customer controlled wavelengths is not solely dependent on the successful deployment of OBGPI

Potential Economic and Competitive Impact of CA*net4 Technology

COOK Report: Isn’t a lot of what you’re doing profoundly threatening to the established Telecom order of things? What can you say about how your agenda is perceived in the Canadian telecommunications community?

St. Arnaud: I don’t think they see this as threatening. In fact many of our telcos are very supportive of these developments. The analogy that we like to use is that we are trying to build, in the optical domain, the equivalent of the first PBX. To date all optical network architectures are equivalent to the Centrex model of networking.

The PBX never killed the telephone companies. And in fact the PBX was quite useful because it helped develop new applications and services outside of what the phone companies could offer at the time. And now gradually you are seeing a lot of companies move away from PBX to Centrex because features and cost savings that drove them to get the PBX in the first place are finally becoming available on Centrex. I think we will see the same cycle with optical networking.

COOK Report: That makes sense. Also I would guess that, given the current marketplace, there are a lot companies out there that would be quite happy to see their fiber get more use.

St. Arnaud: Absolutely. It’s not just the fiber, It’s the wave lengths as well. The fiber at this point is a sunk cost. The issue is much more one of trying to use up the excess bandwidth capacity.

COOK Report: Because in order to light a fiber and get an economic return, they have to make a large investment in optonics that lights some given number of wavelengths no matter how many of those they have or have not sold? Is that the way it goes?

St. Arnaud: Yes. To light your fiber you have to make a very substantial investment in equipment. And therefore telcos become very interested in finding new applications and services that can drive potential customer demand.

That problem started us thinking about customer owned networks, object oriented networking and OBGPI. Today networking is like computing was 40 years ago when the market was dominated by large mainframe computers. But in the 1970s the mini-computer came along followed by the PC which fundamentally changed our thinking of how to do computing. Computing became personal. The user was empowered to develop new applications and services that were not possible on a mainframe computer. With CA*net 4 we hope to move networking in the same direction as computing has gone in the last 30 years.

However, I stress again that OBGPIs experimental. For some unforeseen reason it may not work or have some other deficiencies that preclude its deployment. This will not prevent us from still pursuing our vision of customer owned and controlled wavelengths. But the alternate solutions may not be as elegant.

Start up Issues: Immediate Customer Control of Wavelengths

COOK Report: In such a situation, what could you do?

We can implement a “customer owned wavelength” network with today’s technology. However, the wavelengths initially will be statically configured. If the owner wants to make a change in the routing of the wavelength, it will have to be done manually. Over time we will work with vendors to implement our OBGPI and OON technologies. Our

choice of equipment vendor will largely be decided by which vendor will work with us most closely on implementing these technologies.

COOK Report: So for this early stage of statically-configured, customer-owned wavelengths in CA*net 4 what really changes from CA*net 3? It sounds like the biggest change is that you will have a larger network with many more wavelengths than you do presently?

St. Arnaud: No. From day one we will be assigning ownership and control of individual wavelengths or STS channels to the GigaPOPs, universities and perhaps even individual researchers. They will be free to trade and swap amongst themselves and do whatever they want with those wavelengths. From day one we will also encourage these organizations to directly peer with each other and other international research networks using these wavelengths. But, initially the BGP optical peering will be done manually. Once OBGp is successfully implemented, it will allow these organizations to automatically change the routing of the wavelengths and peering relationships without first contacting CANARIE. So rather than operating a traditional hierarchical IP network as many other research networks do today, CANARIE will only offer an aggregate IP network as an optional service for those organizations that don't need their own wavelengths.

COOK Report: Slide four of your CA*net4 RFI Background Information http://www.CANARIE.ca/advnet/CAnet4_RFI_Background_Information.pdf (an 8 megabyte PDF) assumes CA*net 4 will be operational for ten years. This is longer than the operational lifetimes of the first three CA*nets. What are your considerations behind this long lifespan?

St. Arnaud: The operational life time of CA*net 4 is dependent on government funding. From past experience we have realized that it takes several years to get innovative research applications up and running. Uncertainty as to whether the research network will exist in 3- 5 years acts as a major impediment to the development of new applications. That is why we have proposed a 10 year program for

CA*net 4.

COOK Report: In getting started what are your major milestones?

St. Arnaud: Right now our major focus is on getting a successor network to CA*net 3. We would like very much to be able to purchase a pool of wavelengths. Initially, we may just put simple switches and routers between these wavelengths which CANARIE then will manage. No OBGp in phase one.

Phase two would be to take a subset of these wave lengths and working with some vendors' equipment start experimenting with different approaches for implementing OBGp. It is our hope that we will be able to work with a number of vendors to implement the distributed management of the routing and switching. If funding comes through and all goes well, I would say that by the second half of 2002 we could have OBGp working in some format across the network. It probably won't be a 100% working version and it probably won't be fully robust. It will be an initial pilot version.

COOK Report: If you do get the funding for which you are hoping, you can go out and, within a couple of week's time, buy some wavelengths?

St. Arnaud: That is correct. We have already gone through an RFI process in the anticipation that there will be funding. Consequently, at this point (November 2) we are talking to numerous vendors.

Issues in OBGp Capable Switch Design

COOK Report: When on slide thirty of the RFI Background Information, you say a non blocking STS or gigabit Ethernet or optical switch is possible, what exactly do you mean?

St. Arnaud: An STS switch is a cross point or cross bar. This means that there is no multiplexing across the switching fabric. It sends traffic from one port to another port. Now a gigabit Ethernet switch has a multiplexed fabric. As your packets go in, they get mixed with other

packets that are trying to cross the fabric at the same time. While the aggregate throughput can be higher on the gig E switch, what commonly happens is that you will get what is known as head of line blocking. When you have this kind of fabric, if one port for any reason gets backed up, it clogs the backplane and stops all traffic going to other ports. If you have a common backplane fabric involving shared memory or a shared bus, you are susceptible to the problem and find that a lot of other engineering work must be done to avoid choke points that can lead to its occurrence during the operation of the switch. With routers this is particularly tricky to avoid. You have to employ all sorts of routing policy and make sure your routers are big enough to handle highly variable traffic loads. We would prefer to avoid this.

COOK Report: You also note that an optical switch is possible.

St. Arnaud: Yes. Another type of switch design that can be used in a gig E or in an optical design is cross bar. It has some disadvantages in that, for example, it doesn't allow multi-cast. Think of the capability as putting a physical link or "bar" between two ports. Unlike a router, the paths are "hard wired."

COOK Report: At one point you were going to do OBGp at the routing level. Now you are doing it a layer down with switching? Why? What are the trade offs?

St. Arnaud: Let me answer that question by starting with a quick review of theory. There are two ways of moving traffic through a network. One is with switches, which I will come to in a minute. The other is through routing devices or multiplexers, whether they're classic IP routers, ATM devices or Gig-E devices. Using a multiplexer is desirable when you have a stream of packets going to different destinations. Of course, as the rate of those incoming packets increases, the cost of that multiplexing device increases dramatically. By the time you begin to move from OC 48 to OC 192 speeds, the device becomes exceedingly expensive because the device has to be very, very fast in order to multiplex pack-

ets at those speed without incurring head of line blocking.

The alternative technology is a simple switch. The problem with the switch, however, is that they don't multiplex different streams with different destinations as easily as routers. This is why routers have been predominant in the Internet so far. They give you a lot more flexibility. But switches make a lot of sense if all the packets are going to the same port. If you were to break up a telephone call into packets, all the packets would still be headed to the same port. Thus it makes more sense for telephone calls to use a physical connection between two ports on a switch. But if I'm on a circuit switched line and I want to send some data to someone at the same time as I am talking to you, I cannot do it. I have to terminate our conversation and then call back the switch in order to connect to the other destination. So this is a disadvantage. Nevertheless, since the architecture is connection oriented or point-to-point, it makes for cheap but also high bandwidth capable switches. In this case you just take a cross bar switch and nail it up and let it stay there.

Up to now the trend has been to put larger and larger multiplexed devices at the core of the network. Big routers with OC-192 line cards that can cost upwards of \$200,000 a card because that card has to be fast enough to examine and process each packet in real time. A switch doesn't need to do that, and so it avoids this problem.

Who Controls the Switched Core?

Because of this processing bottleneck, current philosophy says we can't continue to do packet routing at the core. Therefore, increasingly in large IP networks, routing is done at the edge with switches placed at the core. This is not a new concept. Everyone is doing this, even the telcos. But, a question on which views diverge is who controls the switch in the core? Right now, the telcos and most of the telco suppliers are saying that the switched core should be owned and controlled by the telephone company. The customer can have routers at the

edge, but the carrier will go back to its roots and do what it knows best, which is switching circuits across the core. In this case it will do it by setting up an MPLS path for the customer and, of course, charging him accordingly.

We are saying let's try something different and let the customer own the wavelength across the network. This is the big difference. This is why we state that switch architecture for CA*net 4 must allow external users to manage cross connects, provision user VPNs across the switch, and so on.

Optical Switches as Network Objects

COOK Report: Where does your term "object oriented networking" come from? This seems to be new part of your vocabulary.

St. Arnaud: It is a generic term used to describe the use of agents or objects for controlling the switches or the devices. It takes a lot of concepts that we're familiar with in software and applies them in networks instead. This is what Wade Hong has really been doing a lot of work on using tools like Python and SOAP (Simple Object Access Protocol). SOAP is similar to Corba (Common Object Request Broker Architecture).

COOK Report: Please explain further.

Hong: The computing paradigm is now moving toward distributed computing. The idea is that you can have different components of, say, an e-commerce application running at the same time on different CPUs in different boxes while interfacing with a back-end data base server. The question is how do you integrate all this together. Many of the technologies that evolved are based on the Web services paradigm - how to locate distributed elements and how to interact with them.

COOK Report: How does all this all fit into plans for CA*net 4? Does it make sense to ask what kind of "animals" these distributed objects are?

Hong: These distributed objects or net-

work elements are like network agents. For example if we wanted to setup a wavelength from here to Vancouver, an end user or an application would invoke a network agent to set up a light path. This agent would then locate the various network elements necessary to create the light path. It would then communicate with each of the objects or agents controlling the network elements to determine if it is indeed feasible to set up the light path. If so, the reservations would be made and the light path would be created. Similar, a network agent could similarly tear down the light path.

COOK Report: The granularity of these wavelengths is what?

Hong: There are some wavelengths that are 10 gigabits and some that are one. A wavelength is also really just a pipe so you can put what ever technology you want through it.

St. Arnaud: Wade has also talked to some switch vendors who would define a granularity right down through the STS channels to something as small as a DS-3. It depends on the manufacturer and technology being used. However, we are trying to separate ourselves from this and permit the user to be able to decide what granularity he or she wants. It could either be at the level of an entire wavelength, or broken up into smaller components. The granularity would be the user's choice.

When we are talking about agents, objects and Corba and so on, the concept is that we want to be able to separate the hardware from the software. Today software is very much localized and tied to a specific hardware platform. If you are going to separate software from hardware (and this is where Microsoft is going with its .Net strategy), the concept is that you do not maintain any central place for the software. It is a series of links, agents or objects running on different platforms and coordinated to achieve whatever objective you want. Rather than being tied to a hardware platform, what you are doing is linked together through web based objects or agents. We are trying to do the same thing for controlling the network elements which in this case are op-

tical switches and groomers.

COOK Report: But Microsoft is regarded as doing this for network wide purposes of ownership and control. They are going back to the mainframe mentality from edge control. Surely you cannot be accused of this? If not why not?

Hong: You have to examine the objective of why each party is engaged in what they are doing. The objective for Microsoft is totally different from our objective. They may be trying to control the universe, while we are trying to enable our users.

COOK Report: In other words just because you have these architectural concepts, doesn't mean there's only one way that they can be used or applied?

Hong: Yes. For example the underlying technology may be the same but objective behind ".Net" is very different from ours.

COOK Report: Is part of the difference the fact that you can take the ownership of the objects and place them in the hands of the corporations, for example Microsoft, or in the hands of your end users which is what you are setting out to do?

Hong: That's right. A router from Cisco or Juniper today is a big box and you will find that the routing process or routing software is contained within the box. The routing software is very much linked to and integrated with the hardware.

On the other hand, with an object or agent approach, the processes may be located somewhere else. You may run the BGP somewhere else and send the results to the router on a given port. You may, in other words, communicate through an open network object to a forwarding mechanism on the router.

Achieving Hardware Independence for OBGP

COOK Report: But exactly how does

your routing software achieve its hardware independence?

St. Arnaud: This is the problem on which we been working. What Wade has done is to use standard SOAP commands and interfaces to have a routing process or a decision to set up a new path. Our agent enabled OBGP communicates using SOAP with the actual device. That will set up the optical cross connect.

Hong: SOAP is just a messaging protocol that is used to communicate between the objects that model the actual switches and the objects that actually try to control the routes.

COOK Report: It sounds like you're saying that these are processes and that just sit there and run in the device memory and that in the this sense, they are hardware independent?

Hong: That's right. Part of the problem is that if you have all these objects located out there on the network, how do they find out where they are? To answer this dilemma you have directory services as-

sociated with CORBA. This means that you could use a directory services technology like LDAP. With SOAP, we just traditionally used a simple naming service to identify network hosts and locate where objects are to be found. To assist in this there are emerging technologies such as WSDL (Web Services Description Language) and UDDI (Universal Discovery Description identifier). It comes out of the World Wide Web Consortium.

St. Arnaud: Again the principle is to separate the software from the hardware and to have a messaging protocol between these various software objects which could be running anywhere and located completely independent of the routing and switching hardware to which they were addressed.

COOK Report: So are you trying to interconnect the various application infrastructures that have been built on carrier owned wavelengths?

St. Arnaud: In some cases, the wavelengths are provided by carriers and other cases they're provided by the re-

Bill Joy: Era of Agent Based Internet?

Agents Take Center Stage If you attended a couple of these workshops, you would quickly decide that there had to be a better way. At Sun, we have been working toward a better way for about a decade. Our solution is based on agent technology. That is, instead of extending the capability of the network by defining new protocols and having to test the many implementations of the protocols for compatibility, we create the ability to send Java implementations of protocols around the network to machines that include the Java Virtual Machine (Java VM). In this new architecture, the RMI (remote method invocation) protocol by which the Java VM exchanges agents becomes a "protocol to end all protocols."

There are many advantages to provisioning services by agent exchange rather than by defining new protocols: You don't have to write human-language specifications that can be misinterpreted. You don't have to do Connectathons to test compatibility. You can use technology like Jini, which defines services as data types and thereby takes advantage of the Java type system to check the compatibility of the agents' types, which replace types and versions of protocols.

There are other uses for agents in the 21st century, from agents as "shopbots" to agents based on traditional AI technology. But mobile agents as a replacement for hard-to-provision protocols is one of the most practical uses. With the widespread use of Java and the newer Jini technology, mobile agents should become the major way of deploying new services in the years ahead. This change will mark the beginning of the end of the Internet as a world of protocols and the beginning of the era of the agent-based Internet." *Internet Computing*, Vol 4 Number 1 Jan-Feb 2000

gional networks.

I'd just like to back up a minute. To grasp the significance of this new object-oriented paradigm in terms of the Internet and the IETF way of doing things, it is important to consider the paper that Bill Joy, the CTO of SUN wrote for the January - February 2000 issue of *IEEE/ Internet Computing*. In this short piece Joy said that the Internet of protocols is the wave of the past and the Internet of agents is the wave of the future.

The challenge that has faced the IETF and other standards bodies is that as these technologies have evolved to the point where they are dominated by commercial interests. The standards bodies have now become battlegrounds between the big commercial gorillas as each party tries to use standards to enhance its commercial strategy. Because of all the conflicting interests, it has become much more difficult to move a standard through the IETF process

OBGP and the IETF Standards Process

In trying to take OBGP to the IETF standards process, we ran into this problem. Making an open standard in the IETF is very long and difficult process. We also found that if you would get eight different engineers together, you would wind up with twelve different standards for OBGP.

COOK Report: Is part of the problem that those eight different engineers may be employed by of least six different companies with different approaches to switches and routers ?

St. Arnaud: This is part of the issue. We co authored a draft with Viagenie and took the result to the IETF. One of the challenges we ran into right away was that the optical working group in the IETF said: to get this to move forward, you need an endorsement by a major carrier. We said a major carrier is not going to endorse this because they perceive it to be against their own interests.

We will not go away and ignore the IETF,

but the problem is that to implement OBGP within a router we need to get Cisco or Juniper to endorse it and include it within their code. To do all that they first have to see a business case and so on. It could be ten years under those conditions before we see an implementation and an IETF standard. So when Wade pointed out the Bill Joy article and said lets use agents as network objects, I was interested. For if we do it this way, we can implement it and interface with existing BGP processes. With an agent approach we can implement this by using open standards.

COOK Report: But how then do agents replace protocols? If you have these network objects or agents in such a way that you can coordinate their running independently on machines, are you in effect saying that you no longer need a lot of protocols which you would otherwise have to use?

Hong: Java was the first attempt by Sun into implement the ideas that Bill Joy talked about in his *Internet Computing* article. [**Editor:** see <http://www.computer.org/internet/v4n1/joy.htm> and see also text box on preceding page.] SUN has been advocating since the mid 1990s that you should basically have an independent Java agent similar to the agents that we are discussing. The idea is that instead of writing code to control a printer of the network somewhere and have a monolithic application out there that it does it, you should have resources that are highly distributed.

COOK Report: In other words you should be able to query a network node and ask that node to find for you the nearest printing resources that you can use?

Hong: Right. The idea is to put the location of all network resources under the control of the user. Your transport level protocols remain as they were. But you can envision Java as the agent platform on which you can build a lot of capabilities that that otherwise would have required their own protocols. Another tool here is PYTHON that is a scripting language much like PERL.

COOK Report: What you're talking

about seems really different and radical. Is anyone else anywhere to your knowledge looking at this?

St. Arnaud: We are the first to say let's move the ownership and control of the network to the edge even at the routing level. But these object - agent tools are also currently used both with the carrier's centrally managed architecture and with Microsoft. I think that the tools that we are using are all very similar. A wide variety of players are moving toward these object based tools.

COOK Report: But they can be used in many different ways and with different motivations?

St. Arnaud: Well in the telco based world CORBA is a technology that is very hot. But CORBA is very complex. We are using SOAP. Java was proposed as the Internet programming language to facilitate creating agents in the Internet. Architectures such as Jini, Jxta, etc build on this to create an economy for agents and their services. Even in the IEEE there is work going forward on object oriented peer-to-peer protocols.

COOK Report: Where does peer-to-peer come into all of this?

Peer to Peer: an Example from the User Point of View

Hong: The tools that are used to set up peer-to-peer communications are similar to the processes that we are working on for OBGP. Setting up wave length peering is similar to setting up peer-to-peer communication in a NAPSTER network. Instead of loading up music from one server we are saying that we will set up a wavelength to do what users want.

St. Arnaud: Let's assume that you have Napster set up but that rather than pulling down music you want to pull down a petabyte file. If you tried to do this over the commercial Internet with a gigabit Ethernet connection, it would take you about two weeks. In the high energy physics community, they are saying let's take the Napster concept and use it for

our needs. When you click to bring down the petabyte file, the middleware or software underneath will signal the network to set up a wavelength all the way to the source, so when you go “FTP: petabyte-file.name” that will signal the grid to set up a wavelength using OBG. It will use these objects to set up the wavelength to the destination. The huge file would be directly transferred over the equivalent of a virtual circuit set up for that purpose rather than trying to go through a whole bunch of routers.

COOK Report: How will a network like CA*net 4 with all these customer owned wavelengths and all these ‘independent’ circuits fit into the global Internet with all the concerns about the size of routing tables?

Hong: We propose to use cut-through on existing routes to set up optical wavelengths. If you are connected from Vancouver to Toronto you have a normal electronic path between those two points where your routing normally goes. Right?

A user’s application will take the standard electronic routing paths through the CANet* 3 routing tables. The application will then locate available wavelengths and determine whether a light path can be set up. As far as normal routing is concerned, you are still going over normal routes.

St. Arnaud: So in your kernel routing BGP table you will see this new path.

Hong: This path would just be in the table in our router. We are not proposing to propagate that path out to the core routers of the global Internet. The change would only happen within the smaller universe of the routing table of the CANet routers. It would be very rare for these OBG routes to be visible in the larger Internet.

COOK Report: Is what you are talking about analogous to being able to set up and tear down VPNs?

Hong: Yes. There is a very close analogy to that.

St. Arnaud: However, rather than calling up your carrier and asking them to set up a VPN on your behalf, you set it up yourself. The wavelength is just like a VPN.

Hong: You really don’t want to advertise your wavelength to anyone else let alone to the global Internet. You think of what you are doing as cut-through or private peering.

OBGP Proof of Concept

COOK Report: How would you describe what you have done at Carleton University in Ottawa since January 2001? What has developing this as a proof of concept meant? Have you been able to do this in a small regional network where you have gotten institutions throwing wavelengths across a few tens of kilometers?

Hong: We were hoping that would be the case but the Ottawa dark fiber build that we were counting on did not get completed. We are in fact restricted to a single laboratory.

COOK Report: Well a skeptic might say that this sounds like pie in the sky. How would you explain the real meaning of what you have done?

St. Arnaud: Our belief is that research networks should be experimenting with new concepts and protocols that are outside of what is being developed by the commercial sector. We could achieve many of the same objectives of OBG with GMPLS. But given the large commercial research efforts that already exist in GMPLS there is very little additional knowledge a research network can contribute to the field of GMPLS, except perhaps as an early adopter.

By concentrating on architectures and technologies outside of the commercial sphere allows us to make a much greater contribution to the body of knowledge of network engineering. But obviously this entails high degree of risk and possible failure. But that is exactly what research networks should be about.

I have to emphasize that OBG is still

very experimental. This is one of those concepts that may work great in the lab but may not scale or may have some other fundamental flaw when we deploy it in the commercial world.

After a year we may find out that it works great or that it fails because of some fundamental flaw that we haven’t seen yet. It is our job to find this out. We think it does have very significant potential but, hey, we are not perfect.

COOK Report: What have you and Wade found out then with your trial?

Hong: Once we decided on agents we moved ahead and wrote code.

COOK Report: What do you do with your code? Do you, given your lab constraints, simulate what you want to do?

Hong: Emulate would be a better term. We set up and tear down optical cross connects.

COOK Report: You are working with a fairly inexpensive optical switch?

Hong: Yes we are looking at several and working currently with one switch from JDS Uniphase. I am trying to experiment with several nodes so that we can effect switching along a path.

COOK Report: What would you say to the skeptic who complains that you have just done all this in a single lab? What do you need to do for example in the next six months?

Hong: We obviously want to move outside our lab. We were hoping that the Ottawa dark fiber build would be finished by now. Instead it is only just starting.

COOK Report: Why is it late?

Hong: A lot of the delay has been caused by political and organizational issues. Now the Ottawa fiber net aside, for us to actually deploy, means that we have to get a number of these switches. We have only one MEMs module. They are companies with which we have been talking with who are interested in building con-

trollers for these things. The optical cross connect we have is only an 8 by 8 and one problem we have is sourcing a lot of these switches.

We intend through our RFI and anticipated government support that puts money behind the RFI to work, either with an established company or a start-up to get the equipment we need.

COOK Report: Is part of the problem that the MEMs technology you want to use hasn't really been commercialized yet?

Hong: Yes. We have a prototype and don't yet have anything like compact production equipment that fits in a normal switch.

COOK Report: It sounds like you really need a lot more equipment. I didn't realize that you were this equipment starved. In addition to what you just told me, it seems to be very difficult to create this out of nothing because the traditional vendors have difficulty envisioning something that doesn't exist. You have a chicken and egg or "catch 22" situation it seems

St. Arnaud: We are really looking to create a market for this by the build-out of CA*net 4. We believe that this will generate enough interest for some companies to see this as a viable commercial product.

COOK Report: And while looking for the "chicken" to "impregnate" with OBG, from a research point of view, you are being properly cautious?

Hong: While Bill is an enthusiastic advocate, I am indeed cautious by nature. We have indeed been going through an evolutionary process. In the process we have also learned that the object and agent approach is a lot more pragmatic than the protocol approach at this point.

Supporting Object Oriented Protocols

COOK Report: When in your Background Document, you go on to say

"Support for OBG and object oriented protocols encouraged," what do you mean by support?

St. Arnaud: Let's look at two underlying philosophical principles: one stems from the ideas of an economist named Coase. He won a Nobel prize a year ago with a very powerful economic theory that shows that by allowing people to buy and trade among themselves the market is much more efficient than if there is some mechanism of centralized control.

We are saying the same thing for networks. One approach is for the telephone company to decide the optimum path for you to get across the network; our approach is to say that the customers should decide for themselves.

COOK Report: Fair enough. But when you say support, I was more trying to tease out what it meant from the point of view of design of the hardware and software.

St. Arnaud: In response to our RFI a number of vendors have now committed to support some variant of OBG. They are saying they are ready to work with us to adapt some of the software and concepts to their hardware in such away that the OBG agents will run on their hardware. Talking in generic terms, there are a number of different approaches.

I would say that none of the vendors is yet prepared to support the pure vision that we have outlined with the work that Wade **Hong** has done. What they are prepared to do is use something called a proxy. Proxies sort of do what we want. From the user perspective, the difference should not be recognizable. It takes us about half way there.

COOK Report: Are they afraid that if they make a switch that is really dedicated to the implementation of this technology there won't be enough of a market for it?

St. Arnaud: Yes. That is it exactly. Most of the vendors we have talked with are focused, not surprisingly, on the carrier market.

COOK Report: So what precisely does doing it by proxy entail?

St. Arnaud: Doing it with a proxy will necessitate the building of another software shell around their product that emulates what we want done.

COOK Report: Ouch! What does this do to your time line?

St. Arnaud: We aren't worried. We intend first to implement a very simplistic solution using current technology with a few tweaks on it. We will then continue to work with the vendors to augment the equipment. OBG won't be there day one. It will take a process lasting from several months to as long as a couple of years to continually refine, upgrade and add the new software necessary to achieve our goals. Eventually maybe we will even change out the firmware. It is very much a research project.

Network Management System

COOK Report: How are you treating network management issues?

St. Arnaud: One of the problems that we have been struggling with is that there is a real telco mind set among all the vendors. It is very hard to get them to understand some of the principles we are trying to develop here. Consider the traditional Network Management System. It assumes that each carrier is going to manage his network as a cloud into which no one else can see. One of the popular trends is for something called CMN or Customer Network Management views. They have a big management system for themselves and they may give you via the web a little peak into the portion of it that belongs to you. But you cannot yourself directly do anything with it.

The vendors keep coming back to us and asking if we want this customer view capability. We say no. That is not what we want. A customer window must actually directly enable customers to change the network that it shows. While we think that our concept of OBG and customer

owned wavelengths seems straight-forward and simple, we keep being surprised by how difficult it seems for the vendors to be able to grasp it.

We have been saying that the management system needs to be thought of as having four layers. First, there is to be a physical partition layer. One where you partition the switch. The analogy I like to give is that this should function like a carrier hotel where each customer gets its own room in which it does what it wants.

COOK Report: A MIB is?

St. Arnaud: A Management Information Block. Usually one is delegated to each customer. Right now all the MIBs report to the central Network Management System. We want something different. We want customers to be able to direct their own queries to the MIBs without having to go through any kind of central management organization. The MIBs might communicate directly with the customer through a messaging protocol such as CORBA or SOAP.

COOK Report: So presumably current switches are designed with a single MIB for operation by the single carrier to which they have been sold?

St. Arnaud: Yes indeed. This is precisely the problem. We are trying to get the vendors to understand that it is the object-oriented software - Java, Jinni, and SOAP - that really enables our vision.

COOK Report: So you can achieve this opening up and division of the MIBs via the object oriented software?

St. Arnaud: That is correct. Today, almost all the vendors tightly bind the software on their switch to the hardware of the switch. With these object oriented tools we are saying let's decouple the software from the hardware. When you do this you can get routing or switching software that runs anywhere. We think it could be advantageous for them to do this with some of their routing and switch platforms.

For example, today you usually have a

BGP or OSPF router where the software is not only resident on the router but very tightly coupled to it. Under our model the BGP session could be running on a computer a thousand miles away from the router which is in New York City. This is the essence of the thinking behind OBGp. The problem is that if you take a router or switch and try to partition it into multiple BGP sessions, you don't know how many customers you will have. If you have too many, you may run out of local memory and CPU power.

If you say let's separate the two and have the routing sessions run somewhere else, then it becomes very easy to partition the switch and have the ports controlled by different customers and the software that runs the port is on the customer's computer at a different location.

It is important to note this is not a new concept. It is called distributed routing and a number of vendors have tried to implement it in the past. There are two big differences in our model. First, we want to use an open standard messaging protocol to support the distributed routing. Secondly, we want the distributed router to be within the customer's network management and routing domain and not the operator of the distributed routing network.

COOK Report: When you talk about doing cut through routing with OBGp what are you "cutting through?"

St. Arnaud: You are cutting through interior routers and going straight to a border router at the edge of your (AS) or Autonomous System.

COOK Report: Why are you having to use MEMs in your switches?

St. Arnaud: Because they are really needed for cross connects at speeds of OC192 and above. Solid state cross connects and other variants work well at lesser speeds.

COOK Report: With your architecture it looks like OBGpbypasses routers entirely so that you don't need to worry about a switch router interface? Also can you

have switches that use MEMs as well as other technologies?

St. Arnaud: Yes, with our approach to OBGp you bypass routers entirely. But the important thing to note is that BGP routing policy and advertisements are still maintained even though you now have a switched connection. As to the switches: there are companies that are making hybrid switches from a combination of MEMs and solid state technologies.

COOK Report: So how do you handle the OBGp "agents"?

St. Arnaud: The switch providers will handle the agents by installing them with the software on their machines. You can run the routing processes else where in the network. This avoids the problem of the router running out of CPU power if there are too many routing processes on the switch or router at the same time.

If we are presented with a switch that does what we want, we will purchase it. We are talking with some possible vendors who share our vision and are getting very close to implementing it. Therefore, we may not have to do a lot of work in that area.

COOK Report: Then what you need is money to seed a marketplace to create the first batch of customers for this technology and enable the switch vendors to develop what you need?

St. Arnaud: Yes. Once they understand the possibilities we believe that a market will develop for this although we could turn out to be wrong customer empowered architecture. But if we are right, those vendors who work with us should get an early start in an important new market place.

Managing the CA*net 4 Physical Network

COOK Report: What about the management of the physical network?

St. Arnaud: In a sense the ultimate goal is that no central organization will be re-

quired to manage the network.

COOK Report: In the sense of no longer having to have a national manager for a national network?

St. Arnaud: That's right. In the case of CANARIE, for example, the ultimate vision would be for each of our regional networks in each province to have their own switch and be able to trade wavelengths between each other. CANARIE would work with the carriers and equipment vendors to populate the country with these OBGPs, wavelengths and routers, but it would give control over the network to the individual institutions and provincial networks.

Then someone in Alberta can say I have a switch here and multiple wavelengths and I will trade my switch port and this wavelength that goes to Chicago with someone in Seattle who has one that can take me to California. Today a user must signal by means of a UNI (User Network Interface) to the carrier saying "I need an optical wavelength of such and such bandwidth from this location to that location and so on."

COOK Report: But when you have actual cross connects assigned, the user has a pathway from wherever he is in CA*net 4 to wherever he needs to go in CA*net 4. Therefore he doesn't need to do the equivalent of signaling the carrier.

St. Arnaud: That's right. And these cross connects can be incorporated into the user's routing domain. By analogy in the Internet today you have dozens of exchange points in North America and dozens more scattered around the world. If you are an Internet service provider, you have many connections to different Internet exchange points. What we are attempting to do is to extend the concept of these interconnect points or IXs down another level so that they are accessible by individual customers.

Rather than having a few hundred IXs for a small number of big ISPs, we will strive toward the point of having thousands of IXs for individual institutions and users. Ultimately, as an end user,

rather than depending on your ISP to interconnect you to somebody else, the intent is for you to be able to cross connect to that person just like the two of you were actually interconnecting at an exchange point.

Enabling Customer Controlled Peering

Probably the closest analogy to OBGPs something currently running at the MAEs. It is called Peermaker. Today, if you are an ISP and you are at MAE East, you may say: I want to peer with UUNET. Rather than calling UUNET and trying to find the right engineer and so forth, what you do is send an email to Peermaker saying I want to peer with UUNET and do so with this much bandwidth and a bunch of other factors. Peermaker then sends the email to the appropriate person at UUNET and then they reply saying either I accept, or I do not accept or some other qualification. If UUNET accepts the peer request, the message goes back to the ATM switch and ATM switch then makes the necessary cross connect on your behalf.

What we are doing is taking this same principle - except that we are making it much more efficient and using BGP to do our bidding. We use an OBGPs open session or agent to try to set up BGP peering with UUNET. The OBGPs Open Message or agent says "I would like to set up a peering session with you. I have a port and wavelength on this switch. Are you interested in cross connecting to me?" UUNET can accept your request and send you an acknowledgment saying "yes I have a wavelength and a cross connect." It does this by acknowledging the BGP Open message.

COOK Report: What are the length of time issues involved?

St. Arnaud: Generally with BGP you would not want to set anything up for a period of time shorter than a number of days. What we're saying is that ISPs can do this sort of thing today with Peermaker. What we want to achieve is to make Peermaker kinds of abilities available to millions of customers. We will start with

GigaPOPs, then universities and then big researchers.

If I am a big researcher, and I have to transfer a petabyte file by FTP from my lab to CERN that would take me two weeks to do over an OC192 connection. In the grid world there are now are really files this huge. As a researcher I have access to software called Grid FTP. Therefore in the future when I type Grid-FTP CERN and the name of the petabyte file. The grid FTP software will signal OBGPs to set up an OBGPs peering session with a router in CERN. CERN says yes I will accept your request and you can connect to me at such and such a switch at the STAR LIGHT in Chicago. Here is my port number and my wavelength and we both signal the Chicago-based switch to set up the cross connect connection between us. The researchers simply types "Grid FTP file name CERN" and the petabyte file gets a direct path all the way to CERN without a researcher needing to be aware of what is happening in the background.

Doing this uses a standard BGP approach, except that it does so with OBGPs where it says I want OBGPs to set up this direct optical path. In effect all the researcher will need to do is mouse click on the appropriate places in a GUI interface and it will simply happen. We will use RPSL (Routing Policy Specification Language) which is an Internet Routing Registry < <http://www.irr.net> > tool that allows you to manage your routes. It is often used to see who is at an exchange point and what routes and they have.

Hong: Optical peering points for CA*net 4 to connect to European and American research networks are planned in Amsterdam and Chicago. Beyond these two we have no plans for interconnects with non Canadian networks. In the commercial world, many people are setting up VPNs but not yet VPNs based on user controlled light wave exchange. In the commercial world in dealing with BGP, they are trying address quality of service issues and cost issues. These are not the issues that we are dealing with.

COOK Report: But if you succeed in giving your own customers their very own wavelengths, by definition you have solved their quality of service problems.

Hong: Agreed. There are some companies that are doing some interesting work. NetVMG, RouteScience and Sockeye, for example. The latter is a spin off of Akamai, I believe. What they are doing is similar to what we are trying to do with BGP- namely build control overlays on top of existing infrastructure. They are trying to find the best routes in a multi-homed scenario based on a variety of constraints.

St. Arnaud: In the earlier stages of our deployment where wavelengths may not be ubiquitous, this tool would be used to find out whether and how we could manually set up a route to CERN. Eventually, we foresee all this happening in the background as the result simply of the users' interaction with software their workstation.

We expect the winning vendor to partner with CANARIE in developing IRR tools and NMS on the switch. We hope that within six months we will have something to experiment with a on the switch and that in a year or so we may have something that works really well. But let me emphasize again that this is research. It is conceivable that at the end of the day we may conclude that it was a dumb idea. But we hope certainly that ultimately we will say: this is a really neat product and that the switch vendor will then take it out into the commercial world.

Issues of Scaling

COOK Report: When you say you have questions about how well OBGp will scale, could you be more precise as to what you're thinking?

St. Arnaud: OBGp is a part of a broader program that we call Object Oriented Networking (OON). It could turn out to be suitable only for the high end research community. It may not scale to general use in the huge global commercial Internet.

COOK Report: So you are going to see how scales within the International grid community and if it does OK there, you will gradually expand it to more and more institutions connected into various public sector Canadian fiber networks?

St. Arnaud: We intend, as you say, to seed it and expand it over time watching all the while to see if problems develop.

COOK Report: What type of problems might you run into?

St. Arnaud: BGP itself is a rather old and tired protocol. One of the limitations may be in the ability of BGP itself to scale. It is something that we may need to explore. There are many possible "gottchas" that we may not have anticipated and that may be identifiable only with real-world experience.

COOK Report: Perhaps in five to seven years we may have a replacement for BGP?

St. Arnaud: One of the problems in the Internet at now is that it is extremely difficult to introduce new protocol's like IPv6 or a new inter-domain routing protocol because of the existence of a huge legacy infrastructure. Anything you do has to be backwards compatible with all the prior routing software.

Let's imagine that you had a new protocol to replace BGP and that everyone thought it were the way ago. You would have to introduce it slowly and make it backwardly compatible with existing BGP.

Bandwidth Exchanges

COOK Report: You have announced an intention to partner with industry in the development of peer-to-peer bandwidth trading exchanges. This is something that you characterize as an eBay for bandwidth. Would you explain in more detail what you are talking about?

St. Arnaud: Right now there are a lot of bandwidth exchanges out there. It is a booming business. You have Enron, Rate Exchange and Band X. Most of these are

brokered exchanges. It's like buying a house. You as the customer say "I am looking for this kind of bandwidth from here to there" and they act like a real-estate broker in helping you find it. They try to match buyer and seller and when they do, they make a big commission.

The other model is called a market maker model. This is like a New York Stock Exchange where you have market makers. You say I want bandwidth from here to there and they will guarantee you that bandwidth at a certain price - regardless of whether there is a buyer or seller to be found. Like the market makers in the New York Stock Exchange, they make their profit on the difference in the buy and sell prices. They are pledged always to make a market.

COOK Report: Is Enron the only one following the market maker model?

St. Arnaud: Yes. As far as I know, they're the only one following this business model. Now the third potential model is barter. In the past bartering has been very inefficient. It was very hard to identify sellers and buyers which is why brokers came about. Now thanks to the web, there's all sorts of bartering sites like eBay. Rather than a percentage or a broker's fee, E Bay charges to a single small fee for use of the site.

We anticipate that people will continue to buy wavelengths from brokers and from market makers. But we are also hoping to establish a third model where I, as a university, may have a wavelength to Chicago and I will trade that to another university in return for wavelength to New York.

COOK Report: Then you will seed this process by buying the initial quantities of wavelengths to hand out to your provincial networks and universities. Then your hope is that they will figure out how to subdivide which give them and trade amongst themselves and to go out and acquire new wavelengths when necessary.

St. Arnaud: That is the model. They will be assigned their own wavelengths and

have their own switches and can barter, trade and use them as they need to.

Acquisition of Bandwidth

COOK Report: What then are some of the ways open to CANARIE for the purchase of bandwidth?

St. Arnaud: Carriers with transoceanic cables have generally built the equivalent of a condominium fiber network. The transoceanic cables are simply too expensive and have long been that way for a single carrier to take a risk of building. We are saying let's take some of the pricing models of wavelengths on these cables, models that have been in operational a long time, and apply them to the purchase of domestic wavelengths.

Doing this would enable a customer to purchase from a fiber owner 10,000 km of wavelengths for a period of five years at one negotiated dollar amount. The customer would be able to slice and dice the 10,000 km into the links that it needs and would be able to apportion them between the various physical points of network infrastructure. The distances traveled by each wavelength would be added together to get a total distance for all the wavelengths. If that distance were less than 10,000 km, the total price would be the primary negotiated figure. If it were more than 10,000 km, the excess would be purchased for a lesser amount.

Another way of doing it would be for the customer to purchase what is known as a fiber or equity IRU. Here the customer would own a fixed percentage of the fiber capacity. In other words, regardless of the number of wavelengths established by the owner over that fiber, the customer would get the same percentage of wavelengths. A third possibility is that the customer would pay an extremely low price for wavelengths that a fiber owner had lit but had not been able to sell commercially. With this option, should the owner be able to make commercial sale of a wavelength that CANARIE is using. Then CANARIE is able to use one of the other options to acquire a replacement wavelength. With a wavelength equity

the fiber is lit, and you are buying a percentage of the lit wavelengths. With a fiber equity you buy a percentage of the fiber and you both decide how to light the fiber.

For CA*net 4, we anticipate acquiring access to wavelengths ranging in bandwidth from OC3 to OC192. However, because gigabit Ethernet as a transmission medium to the end user is relatively inexpensive, we are translating OC-"x" STS wavelengths with STS to Gigabit Ethernet converters placed at every GigaPOP.

I want to make an important point here when we say wavelength, we have a much broader definition than a traditional analog wavelength. A wavelength, by our definition may also be a point-to-point lightpath that may be a portion of an analog wavelength — for example a STS channel.

COOK Report: Is this converter a physical device?

St. Arnaud: Yes. A box capable of converting an OC-192 to gigabit Ethernet costs about \$30,000.

COOK Report: And someday soon you should be able to just buy straight ten gigabit Ethernet?

St. Arnaud: That is certainly our hope. In this market how soon is another matter.

Topology of the New Network

COOK Report: If you look at all this from a really high level, do you begin to see the future of the internet more and more in terms of loose confederations of networks? Are these federations combinations of both groups of people and of technologies associated with such groups who, to varying extents, choose to communicate with each other.

Hong: You could view it in such a manner. There are various groups working in parallel to come up with various alternatives. Given the Internet's open end-to-end architecture, they can all try their ap-

proaches.

COOK Report: Let's turn then to the network itself. You say that in the "old world" you have one national carrier cloud from which every one gets a "managed service" while in the new world each GigaPOP operates its own separate network cloud. Slide 20 in your RFI Background Information explains the basic topology of your Customer Empowered Network. It is Figure One on the next page. Would you elaborate please?

St. Arnaud: What is happening already today even without CA*net4 is that many institutions are getting their own dark fiber and are buying their own wavelengths. We are trying to build some infrastructure to enable them to expand upon this process. As a result, when we are able to provide wavelengths to GigaPOPs, then each GigaPOP can begin to act as a carrier for that metropolitan area. In turn, the GigaPOP may hand over control of the wavelengths to universities so that by connecting to each other with OBGp through the GigaPOPs or through other exchange points they can begin to cooperate in their own national network cloud. In this architecture the national CANARIE network's switches are mini-IXs where clouds touch.

COOK Report: So even though it is not labeled as such, the switch in the very center of Figure One is a CANARIE switch?

St. Arnaud: Yes but the ports on that switch are controlled by the different end users.

COOK Report: What about the switches in the City B and City C clouds that are directly connected?

St. Arnaud: They are all the same. For example a university in City A may own a wavelength that goes through the switch in the center and from there over to City B and it may own a second wavelength going to a switch in City C it controls the ports on these switches.

COOK Report: As you start up, your universities or other entities will be routing wavelengths a lot of the time through

No Central Managed Wavelength Cloud

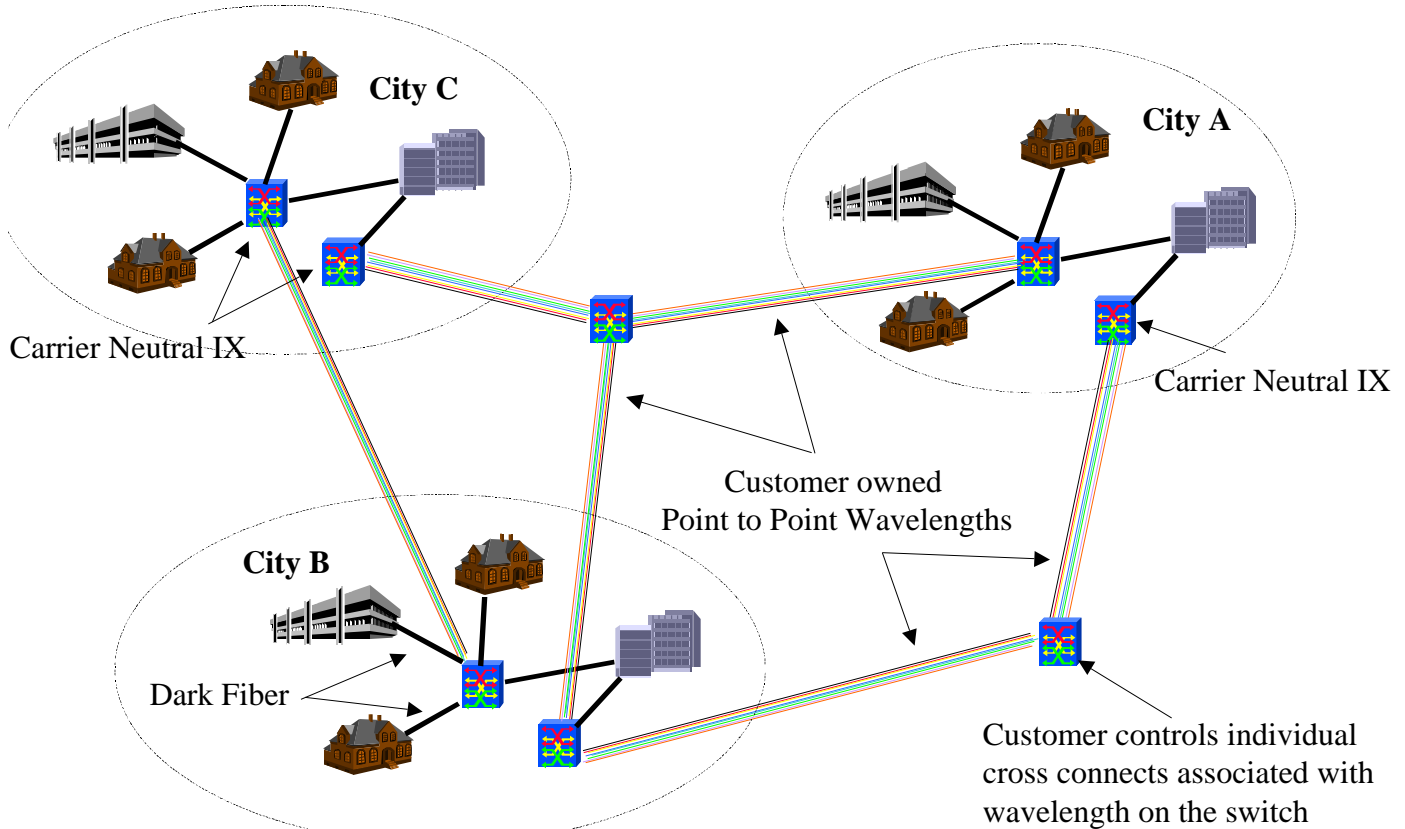


Figure 1: Customer Empowered Network

CANARIE switches. But, as CA*net 4 gets going management control over the wavelengths will migrate and become more and more decentralized?

St. Arnaud: Correct. For example, let's say the regional network in British Columbia (BCnet) is given control over four pair of national wavelengths from CANARIE. BCnet then transfers control over of two pair of the wavelengths to the University of Bristish Columbia which in turn gives one pair to TRIUMF, the high energy physics research facility at the University. TRIUMF now has its own national network cloud with a pair of wavelengths that can run on each CA*net 4 switch. Initially, TRIUMF may connect to SNO, another physics research facility, at Sudbury. A bit later it could decide to switch its wavelengths to the STAR LIGHT to connect to Argonne Labs in Chicago. Or it could chose to drop off a portion of its wavelength as an STS channel at Calgary and another in Sudbury. All the time the topology and routing of the wavelength is under control of the researchers at TRIUMF. The decentralization we envision will give our cus-

tomers extraordinary flexibility in what they do with their resources.

COOK Report: Looking further at the RFI documents, what is an ORAN?

St. Arnaud: An ORAN is an Optical Regional Advanced Network.

COOK Report: Like the Alberta provincial optical network I wrote about a few months ago?

St. Arnaud: Right.

COOK Report: Do they have those in the various stages of development in just about every province now?

St. Arnaud: Every province has or will soon have an optical regional network. Many of them are based on dark fiber optical networks. In British Columbia, for example, they have in an optical dark fiber network linking hospitals, universities and some schools. Alberta has the SuperNet project of course. Saskatchewan and Manitoba have wavelengths and Gigabit Ethernet services

from a carrier. Ontario is building out a dark fiber network. RISQ in Quebec has already done that. Also, the Maritime Provinces are in the process of putting together a plan for a fiber network.

COOK Report: Are they going to be condominium based like Quebec?

St. Arnaud: That is what they hope. Until they get the responses to their RFP and their funding taken care of, it is an open question.

COOK Report: If you look at what's happening in your provinces with these networks, are there any commonalities in what the provinces are doing?

St. Arnaud: It varies from province to province. But I would say there is a universal acceptance in the value proposition of "owning" a network as opposed to leasing. Ownership can be through dark fiber or through wavelengths. Dark fiber, however gives you the most flexibility in terms of network architecture.

COOK Report: Let's turn to slide 24 Pos-

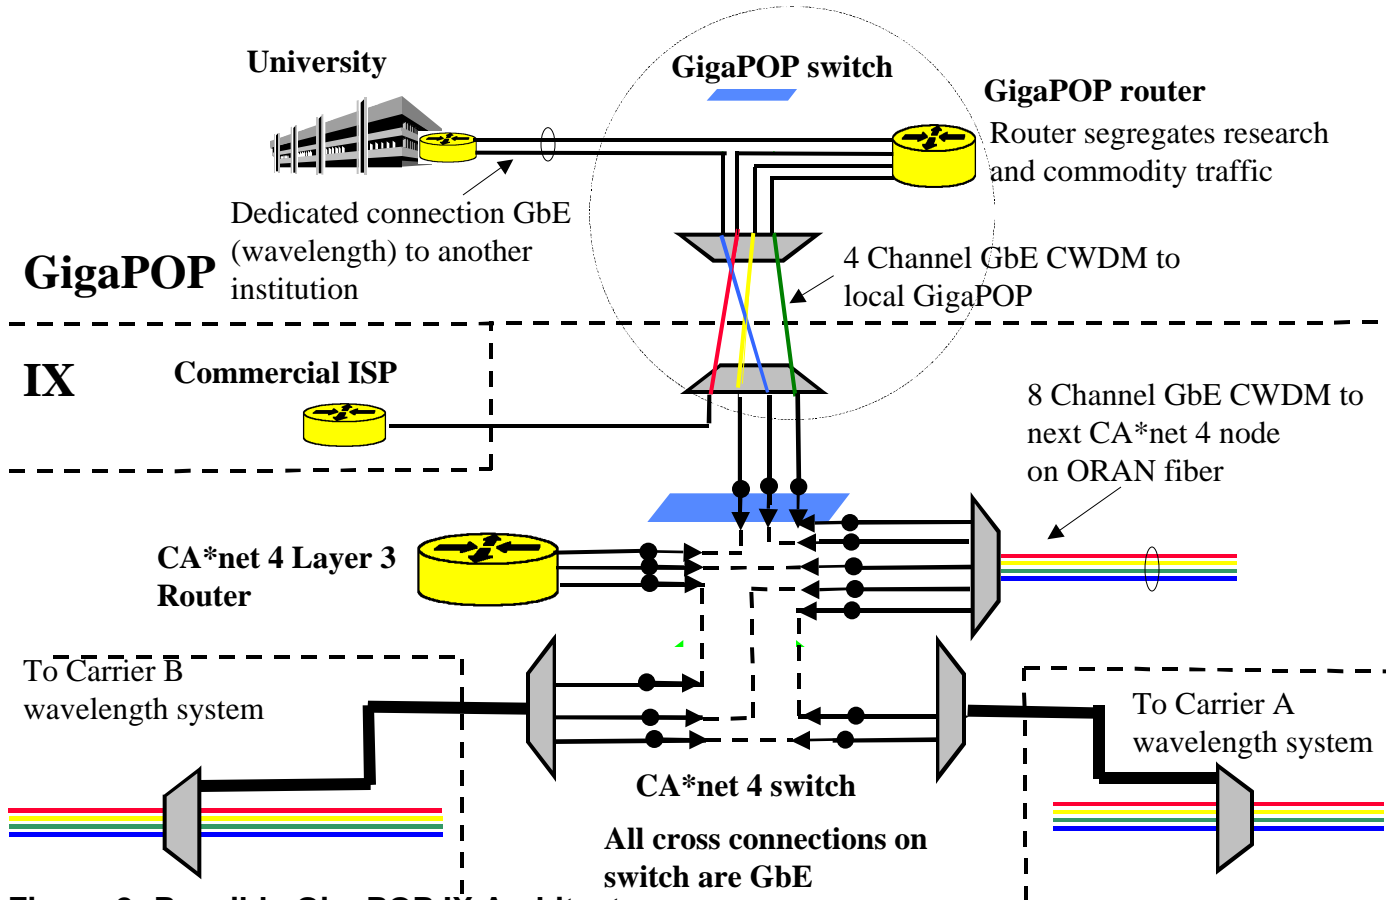


Figure 2: Possible GigaPOP IX Architecture

sible GigaPOP - IX architecture shown above as Figure Two. What is going on here?

St. Arnaud: The GigaPOP not only connects to CA*net 4 but also to a commercial Internet service provider. It's function is to extend CANARIE services and commercial services to the local Ca*net 4 customer community.

COOK Report: At the bottom your CA*net 4 Layer 3 Router and CA*net 4 Switch show your national bandwidth while the circle around the GigaPOP Switch shows the interconnects at the local university GigaPOP level?

St. Arnaud: Correct.

COOK Report: OK. Then you go on in slide 27 Wavelength Assignment and Figure Three on the next page to show a different view. However it is one that just like the GigaPOP architecture shows national bandwidth from two different carriers. A and B. What's happening?

St. Arnaud: This reflects one of our

challenges. We may have ultimately 200 universities and research centers that we wish to serve. However, we may have only 50 wavelengths to work with. The question therefore becomes how do we allocate and assign them?

Here you may begin to get as much into politics as into technology. But, as a first priority, we want to provide wavelengths to those institutions and researchers who have need of a direct, end- to- end, high bandwidth connection. This slide shows how we might take wavelengths from the two carriers and then give control over the wavelengths to regional networks which in turn hand control over a subset of the wavelengths to key universities or to researchers who need the end- to- end bandwidth connection.

The goal is to get a pool of wavelengths and then divide that pool among the ten regional networks based on demonstrated need for a high bandwidth connection. Let's say a regional network gets five wavelengths. It in turn may divide the wavelengths among the institutions they need to serve.

Slide 27 Wavelength Logical mapping (Figure Four on the next page) shows how the universities in AS6 and AS1 might use CANARIE switch number 9 in carrier B's cloud to connect OBGp peering sessions.

Figure Four also shows how the regional network (ORAN) at AS5 might use an OBGp link to peer with regional network (ORAN) at AS 2 via ports on CANARIE switch 10 in the cloud belonging to Carrier A. With each wavelength that a university gets it also receives a set of ports across the CANARIE switches to use in linking those wavelengths in ways that best suit its needs. The university can then do whatever it wants with the bandwidth.

Slide 29 (Figure Five on page 17 below) Customer View of Networks shows how, if you are a network administrator at one of those universities, your network looks different now. Rather than having a campus LAN to administer you have a LAN that extends or can extend across the country because of the national wave-

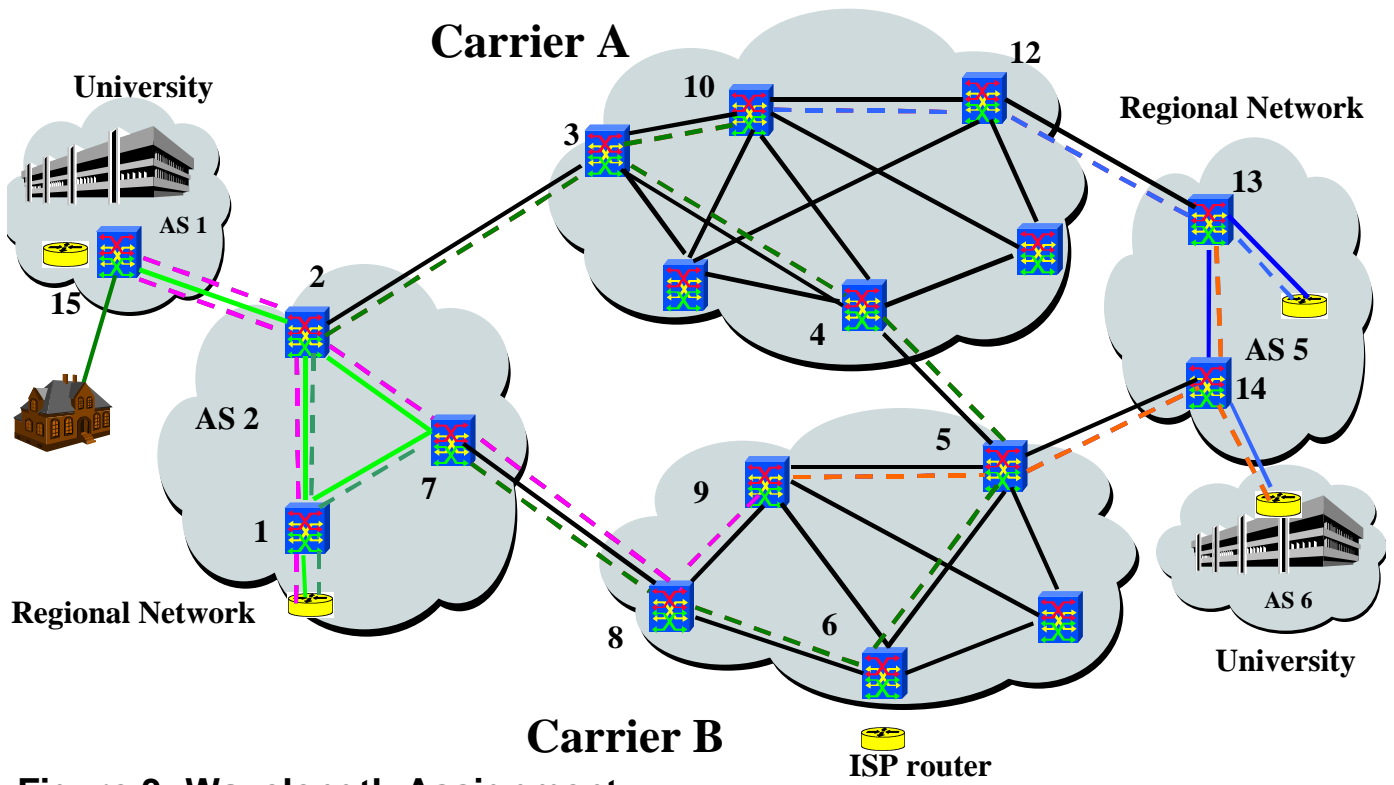


Figure 3: Wavelength Assignment


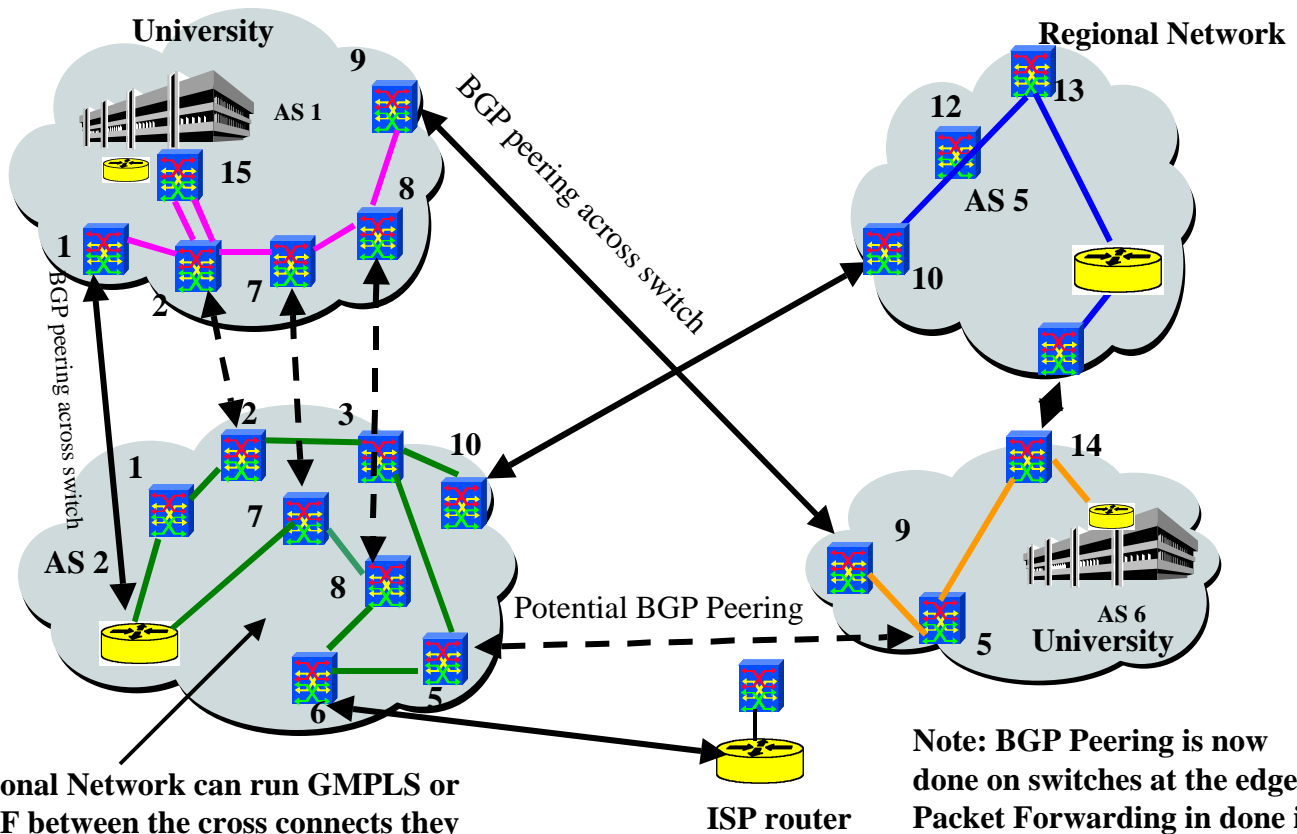


Figure 4: Wavelength Logical Mapping



Regional Network can run GMPLS or OSPF between the cross connects they own on the switches
Figure 5: Customer View of Networks

lengths that you now control.

COOK Report: Finally your Slide 22 Wavelength Scenarios (Figure Six on page 18 below) is useful to look at because, in addition to showing the basic geographical topology of CA*net 4, it shows the three levels at which OBGp will function?

St. Arnaud: Yes. You have GigaPOP to GigPOP. For example BCnet to RISQ. University to university: for example Vancouver to St John's and workstation to workstation. Each of three different wavelengths is shown traversing through eight CA*net 4 switches.

A New Foundation for the Future of the Internet?

COOK Report: It seems that the Internet is continually in search of a business model. Perhaps by 1996 or 1997 we found a viable business model for dial up and leased lines. But by 1998 DWDM

and optical gear were making their impact felt. Over the next two years they threw all earlier assumptions out of whack. Could it be that the Internet and telecommunications in general now needs a new business model that takes the optical revolution into account? While the problem of the local loop still looks especially nasty, Ethernet in the first mile development looks very interesting. Also there are new ideas emerging about how to get fiber to the home. See for example: http://lw.pennnet.com/Articles/Article_Display.cfm?Section=Articles&Subsection=Display&ARTICLE_ID=0252

What are your thoughts on what this model might look like and how it might work itself out over the next five years?

St. Arnaud: I think the business and architecture model of the future will be of control and management moving increasingly closer to the edge, not only of the in terms of applications, but also in terms of control over the infrastructure. I think that one of the drivers for this will be as

a consequence of the issues that of Larry Lessig has raised where content and distribution companies are trying to exert control over the Internet infrastructure to protect their intellectual property interests. Decentralization and minimizing control at the center will help thwart these challenges. We are already working on concepts with our industry and research partners to extend this concept of customer control of wavelengths all the way to the individual home.

For example, one of the concepts we are exploring for delivering broadband to the home is where the homeowner owns and controls individual strands of fiber. We are investigating how to extend the principles of OBGp and distributed computing to the consumer level through concepts such as Reverse Passive Optical Networks (RPONs). With RPONs the control of the wavelength and fiber routing, including moves, adds changes is always under the control of the end user. From time to time the end user may delegate that control to the service provider of their choice, but the fundamental prin-

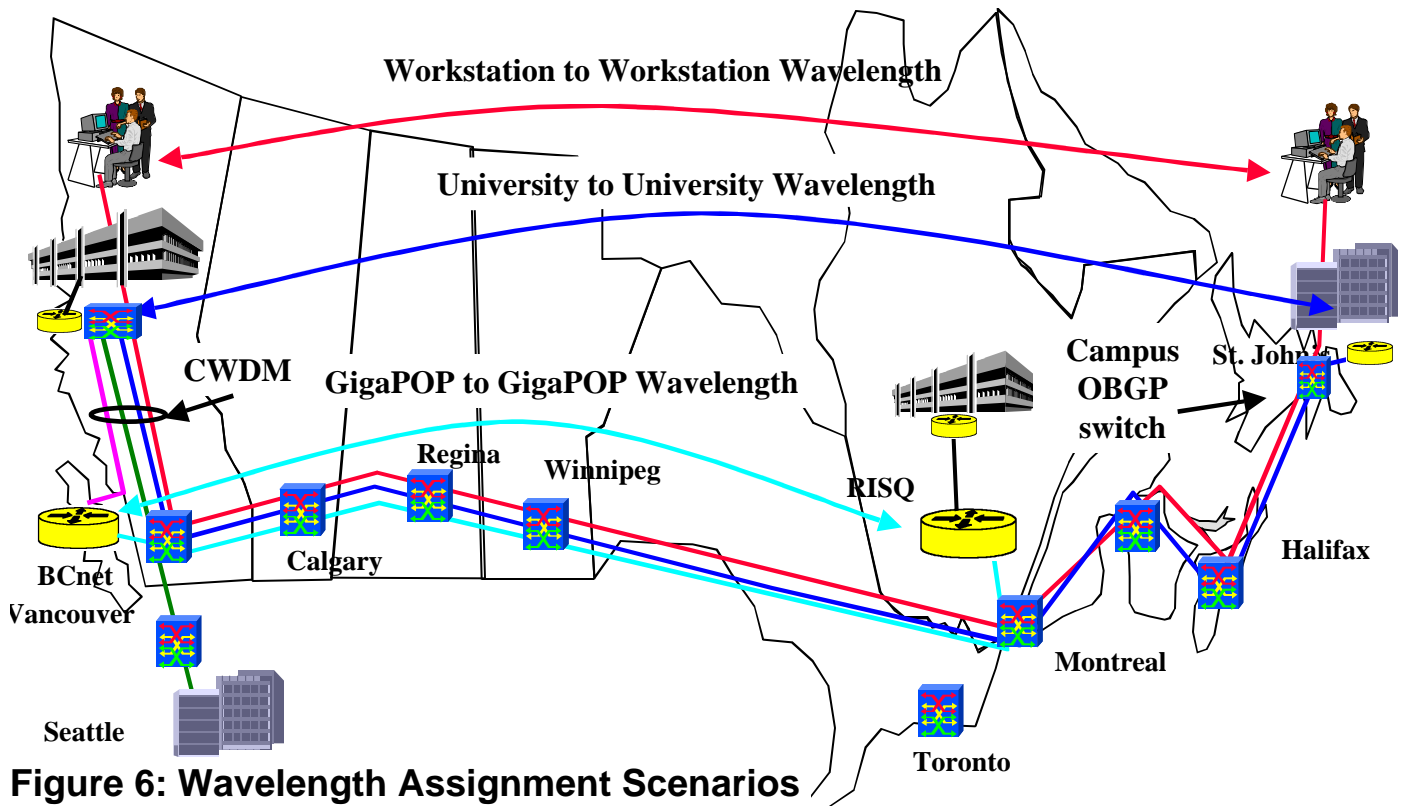


Figure 6: Wavelength Assignment Scenarios

principle remains the same - end user ownership and control at the edge.

In the future we see a physical network infrastructure that closely parallels Morpheus and other peer-to-peer networking paradigms. The end user will have a choice of whether they wish to subscribe any number of "walled garden" service providers. Or they may choose to physically connect to community networks to share files and data with high speed Gigabit wavelengths bypassing all traditional hierarchical service providers.

A lot of these concepts are still very speculative and unproven at this point in time. But what it does point out - is the critical role that research and education networks still play in the ongoing development of the Internet. The Internet would have probably never been developed if we restricted ourselves to the telecommunication business models of the mid 60s and 70s. The Internet only came about because universities and research centers began to explore new ways of interconnecting computers. In the late 1990s the conventional wisdom was that university research networks had done their job - the Internet innovation phase was over and it was time to

focus on commercialization and scaling issues. But in reality the Internet revolution is far from over. Most importantly research and education networks will continue to play an absolutely critical role in exploring new concepts in networking which initially may appear very radical to the traditional telecom world.

Many of the ideas we have talked about here may fail. Many may never reach commercialization. But that is the critical role for research and education networks - to explore ideas and concepts that are orthogonal to main stream thinking and carry high degree of risk. That is how the Internet started. And that is how I believe the next revolution in the Internet will be started.

Editor's Afterword

The Canadians are taking very seriously what the United States no longer does: (1) the assumption that publicly supported infrastructure is important. (2) that government pump priming for such infrastructure will support innovation which in turn by promoting economic growth will more than pay back government investment. The idea of a commons and a national public interest is

alive in Canada as evidenced both in CA*net 4 and the Report of the National Broadband Task Force. http://broadband.gc.ca/Broadband-document/english/table_content.htm

Key objectives of the Broadband recommendations are (1) to ensure third-party open access; and (2) to ensure competitive and technological neutrality by means of condominium fiber builds to all communities and carrier neutral hotels in every community. Public funds are expected to be available as infrastructure subsidies only to companies that agree to build and maintain networks rigorously open to third party access. It is anticipated that broadband in Canada along with CA*net 4 will receive budgetary support during the second week of December. Canadian's are well aware of the arguments made by Larry Lessig's new book as to the importance of the maintenance of a commons and open access to that commons. [See review of Lessig's book pages 25-30 below.] Meanwhile Ca*net 4 is a completely separate research network - a walled garden with lots of entrances. It seems likely that lessons learned in CA*net 4 will spill over into the grassroots broadband infrastructure that Canada is preparing to build. If they do, the face of global telecom will be radically changed.

Trust as Qualified Reliance on Information

Ed Gerck
egerck@nma.com
Summary

*"If the world were really random, chemistry, cooking, and credit would not be possible, so our models cannot be figments of our imagination."
P. Cheeseman, "Finding the Most Probable Model," p. 91, 1990.*

Editor's Introduction

Trust is a word that is commonly applied to many situations and consequently has many shades of meaning. The following essay by Ed Gerck focuses on one precise set of coherent meanings: the concept of trust in the context of communication. More specifically, in the context of the engineering problem of Internet communications. At the same time he demonstrates why trust is needed in this context. Trust is considered something essentially communicable, but with specific rules for its communication. Gerck's exposition also discusses the induction (communication) of trust in heterogeneous environments, from human to machine, machine to machine, and machine to human.

By allowing trust to bridge the many gaps between human and machine, people will be able to tailor their own human to human communication needs via the Internet. What this means is that communities of interest, as networks of people, can build their own networks within the Internet according to their needs, without any limitation imposed on them by their Internet connectivity.

No one is at the edges of the Internet-network, while everyone is at the center of their own network. In this sense the flat, edge controlled Internet that we wrote about in the lead article in the December 2001 *COOK Report* is really just a local vision of a multi-dimensional network of networks made up of many different user groups and their networks, who actually act as control centers of such local networks.

Einar Stefferud observes:

I have known Ed Gerck since 1997. I have discussed with him and read many of his previous papers on trust; so this essay serves to bring together many different threads that we have discussed on- and off-line since 1997. So I now see that all my previous talk about the Net being edge controlled needs to be revised in some new framework. In short, the Internet does not really have a center or edges. It only has connection points, each of which can be connected to any other such connection point for the purpose of packet exchange. One reason that the Internet does not have an edge (as I just realized) is that at any termination connector, it is possible to extend the Net beyond that point by relaying packets, or by relaying messages, via dial-up modem, FAX or printer, or word of mouth, for that matter. So we suddenly discover that we cannot define any edge of the Internet.

Gerck's communication concept of trust may be just the beginning of a broad-based understanding of a new view of the Internet where security is a core part of the design. This new view of the Internet is that of a large collection of local network centers and edges, of potentially overlapping subsets of the total Internet. It is built around local common interests and purposes (communities of interest), but with a global communication pattern that closely resembles how we humans communicate across such boundaries and how our commerce works; and it looks like an assembled collection of networks, each of which has its own local centers and edges when we observe them closely, but the global collection of these local networks into the whole Internet doesn't have a single global center or any edges.

Now, since we have so many available connections, Gerck is saying, let's use

sets of connections to enable us to transfer trust using distinctly separate multiple channels. Except that, in the essay that follows, he leaves for a next article the discussion about how one can use those multiple channels to induce trust, and how many channels to use. First, one needs to establish the need to use multiple channels, before explaining how to use them.

The problem is that if the Internet is this thing with users and servers attached to its interconnection spigots with nothing but connection pipes between them, and where any attached user or system has protocol-based communication access to all others so mounted, then we must ask what controls the whole thing? And where might we mount a controller for the whole Internet?

The essay explains that the answer depends on how you use the communication concept of trust. You may choose to be at an edge of some local network by joining a mailing list or participating in a message board on some website, or subscribing to some information services, trusting that which has been authorized for you. Or you may choose to be at the center of your own network where you control the nature of all the connections. Or you may choose to be at both, edges and center, and from this position you will be able to realize the full potential of the Internet. However, quite independently of your choices, the Internet is still just a Network of Networks—which prevents anyone, from anywhere, from hosting a single control center for all elements of all networks in the Internet. We do not even know what local networks exist inside the Internet.

Trust as Qualified Reliance on Information

by Ed Gerck

When I say that the key by which to solve the fundamental problem of Internet communication is trust, I usually get two reactions. The first is "what is the fundamental problem of Internet communication?" The second is "what is trust?"

Let's answer these questions.

In 1948, Claude E. Shannon created information theory and stated that the fundamental problem of communication is that of reproducing at one point either exactly or approximately a message selected at another point. Frequently the messages have meaning; that is they refer to or are correlated according to some system with certain physical or conceptual entities. These semantic aspects of communication are irrelevant to the engineering problem. The significant aspect is that the actual message is one selected from a set of possible messages. The system must be designed to operate for each possible selection. [quoted from Shannon, C. "A Mathematical Theory of Communication," *Bell Syst. Tech. J.*, vol. 27, pp. 379-423, July 1948. Available at <http://cm.bell-labs.com/cm/ms/what/shannonday/paper.html>]

Fast forward to 2001. The fundamental problem of Internet communications is that of reproducing at one point exactly the message selected at another point. No one can control at the same time both ends of a connection to another party, neither sending nor receiving. The route followed by the messages cannot be controlled or positively verified by any party. All messages have meaning relevant to the engineering problem of transferring the bits, but with different meanings, some of them intertwined, at different protocol layers; that is the messages contain packets, headers and fields that need to be formatted, interpreted, or verified, in correlation with references, by suitable protocols at each end and en route.

Reading the two paragraphs above, we realize that Shannon's information theory fails to model Internet communication. Internet protocols are much closer to human communication than to Shannon's idealized communication systems. We, speakers of the same natural language, communicate with one another by trading contents, not by exchanging uninterpreted strings of symbols (bits). Each bit of information sent by a human to another must either contribute to the content or be discarded. Content may have different meanings, some of them intertwined, at different layers of our understanding. And, in the same way that content must be conveyed in human-to-human (H2H) communication, we find that content must also be conveyed at different protocol layers in Internet machine-to-machine (IM2M) communication – not just bits. In internetworking, machines are not just trading uninterpreted strings of symbols, or bits. They are trading bits and meaning, machine to machine. They are talking.

Before anyone thinks that I intend to turn Internet engineering into some sort of socio-philosophical-techno babble, let me comment that the objective here is to discuss a technical solution to the engineering problem. We are still happy with Shannon's definition of information as a measure of the decrease of uncertainty at a receiver. In other words, information is what you do not expect. However, the problem has now an added dimension. We must be able to convey meaning in IM2M communications. But this meaning is not the same meaning conveyed in the H2H communications using those same machines, and not the same meaning at every protocol layer either. Meaning must be conveyed in heterogeneous environments, from human to machine, machine to machine, and machine to human.

Introducing meaning into information theory, so that the communication of meaning can be described, has been an open problem since 1948. I assert that the way to communicate meaning is to first communicate trust and bits, and then use them to define the meaning. This may

sound like jumping from a frying pan into the fire, because we must still communicate trust. However, my assertion is based on the observation that trust is essential for H2H communication and needs to come first before we can rely on the contents being communicated. Thus, since we can readily observe that communication processes in general are in many ways very close to H2H communication, as exemplified by the IM2M communication discussed above, we should expect that trust may also be essential for communication in general and also needs to come first. In other words, rather than introducing meaning into information theory, we introduce trust. Meaning will be introduced and conveyed implicitly.

But what would trust be in the context of IM2M communication? Or in the context of communication processes in general? It would need to represent the same abstract idea of trust in the context of H2H communication. What is this abstract idea?

The answer to these questions must be useful for a wide range of communication systems, such as H2H, IM2M and others that need to interoperate.

The only answer that turns out to be viable is that trust in communication systems must have nothing to do with feelings, emotions or other psychological and multiple-variable concepts. Trust is to be understood as something potentially communicable.

Further, trust must bridge different instances and observers, otherwise communication would be isolated in domains with islands of user interoperability that could not be bridged over time. This means that different subjective, objective and intersubjective (see Glossary) realizations of trust must depend on some common, basic and abstract expression. This expression is simply:

"Trust is that which is essential to a communication channel but cannot be transferred from a source to a destination using that channel."

Additional reason to use this expression as the definition of trust in communication systems and why it is useful for Internet security protocols is given in the sidebar, quoted from the book *Digital Certificates: Applied Internet Security* by Jalal. Feghhi, Jalil. Feghhi and Peter Williams.

It is important to note that there are also “poetic” or “everyday” uses of the word trust that permeate some security work and Internet communication protocols. This may explain trust’s “bad name” as a difficult concept and as an overloaded terminology. The problem is not however in the concept of trust by itself, but in using trust quantifiers that are either artificial or limited in the context of communication systems. A common limitation is to use trust as a synonym for authorization. However, this is valid only in a network, not in an internet. For example, in a network a trusted user is a user authorized by the network management to access some resources. But where is the “Internet management” in the Internet? It does not exist.

Let’s use the definition of trust just given and move toward an understanding of it by using simple examples and problems that can later on be translated to more complex H2H and IM2M communication.

Examples and Problems

To make progress in understanding all this, we probably need to begin with simplified (oversimplified?) models and ignore the critics’ tirade that the real world is more complex. The real world is always more complex, which has the advantage that we shan’t run out of work. [quoted from Ball, J. "Memos as Replicators," *Ethology and Sociobiology*, vol. 5, p.159, 1984]

Let us suppose that a lion sees a lamb and tells the lamb “I’m not hungry.” What should the lamb do? Usually, the lamb would run away – but not necessarily. There would be no danger if the lamb were able to know with a high level of reliance, acceptable to the lamb,

Trust Points

Commencing with a quotation from [SHAN48], Egardo Gerck leads an Internet discussion <http://www.mcg.org.br/trustdef.txt> with the assertion, used with permission, that:

"In Information Theory, information has nothing to do with knowledge or meaning. In the context of Information Theory, information is simply that which is transferred from a source to a destination, using a communication channel. If, before transmission, the information is available at the destination, then the transfer is zero. Information received by a party is that what the party does not expect—as measured by the uncertainty of the party as to what the message will be.

Shannon’s contribution here goes far beyond the definition (and derived mathematical consequences) that “information is what you do not expect.” His zeroth-contribution (so to say, in my counting) was to actually recognize that unless he would arrive at a real-world model of information as used in the electronic world, no logically useful information model could be set forth!

Now, in the Internet world, we have come to a stand off: either we develop a real-world model of trust or we cannot continue to deal with limited and fault-ridden trust models, as the Internet expands from a parochial to a planetary network for e-commerce, EDI, communication, etc.

And, what would be this “real-world model of trust” for the Internet world? Here, akin to Information Theory, trust has nothing to do with friendship, acquaintances, employee-employer relationships, loyalty, clearance, betrayal and other hard to define concepts.

In the concept of Generalized Certification Theory (see <http://www.mcg.org.br/cie.htm>), trust is simply “that which is essential to a communication channel but which cannot be transferred from a source to a destination using that channel.”

Dr. Gerck’s underlying observation that the integrity of certificate-based security systems (such as the X.509 Authentication Framework) hinges upon the very notion of trust is very valuable. Although trust is defined by the X.509 standard in terms of integrity—a CA is expected to reliably perform its user authentication and certificate registration duties—this does little to establish any conceptual properties of trust itself as a basis for building secure systems.

As anticipated by the ISO process, local environments are expected to profile and tailor the X.509 Authentication Framework. In this way, they can address the integrity requirements of national, application, community, or personal needs. ISO data communications standards are generally constructed assuming that islands of user interoperability will form, and the economic or social benefits of networking will inevitably cause systems to link together over time. X.509 does not impose a particular economic or social model of integrity, however. Such telecommunications standards generally limit their scope to stating technical matters. Models of integrity and trust in a particular space are best left to the communities of interest, governments, and industry forums, which are most familiar with these groups’ specific needs.

Excerpted from *Digital Certificates: Applied Internet Security* by Jalil Feghhi, Jalal Feghhi and Peter Williams, Addison-Wesley, ISBN 0-20-130980-7, p. 194-195, 1998. Copyright work.

that the lion is not hungry. Can the lamb trust that the lion is not hungry?

Let us follow some idealized steps of this communication protocol.

The lion does not need to receive any information from the lamb besides that which is communicated in the communication channel itself – the lamb is there. The lamb obviously needs to know whether the lion’s assertion “I’m not hungry” is truthful. The truthfulness of this assertion is not information and cannot be transferred using that same channel.

Why not? The truthfulness of the lion’s assertion cannot be information because in Information Theory information has nothing to do with knowledge or meaning, and it cannot be transferred using that channel because how could the lamb know that the lion was not lying?

The same situation would apply to two machines on the Net, or to humans communicating. Information is surprise, and not always nice. Therefore, we see that “trust me” is an empty affirmation; a self-affirmation cannot communicate trust.

In other words, a decision to trust someone, the source of a communication, the name on a certificate, or a record must be based on factors outside the assertion of trustworthiness that the entity makes for itself.

Loosely speaking, we can say that “information is what you do not expect” and “trust is what you know.” In the lamb example, the lamb needs to trust whether the lion is hungry. This could not have been information, because information is what the lamb does not expect – information is surprise. To be sure, the lamb does not want surprises in regard to the lion’s appetite.

Linking both concepts of information and trust, we can say that “trust is qualified reliance on received information.” This definition, derived from the first definition, is just one out of dozens of other trust definitions that can be de-

rived, all coherent with the first definition.

This example also shows the interplay between trust and power. A very large difference in power, of one agent over another, implies that the more powerful agent can offset and control the other agent to such a degree that the other agent’s actions are immaterial, even if the actions are already occurring – hence, a vastly powerful agent does not need to trust the least powerful agent. On the other hand, the least powerful agent needs trust in the other agent’s behavior, since it cannot offset or control the other agent’s actions to any degree – it needs to know with high reliance what the other agent’s actions can be and, in some cases, what they cannot be, before they happen.

Let’s look at some problems. How can I trust whether the message received by a party is the message that I sent, if I cannot control both sides of the communication channel? This question has gained in importance lately. Over time, we are finding that everything we see on our screens just might be false, including e-mail that says it was mailed by friends, or even digitally signed by them. And we are never totally sure that the website pages we are looking at are really from where they say they are from, or that what they say was not tampered with. Isolated networks do not help. Why? Because we are also finding that authorized users moving data from exterior to interior networks can compromise purportedly secure networks even if the networks are fully isolated.

These problems have no solution today. And it is well known that digital certificates and cryptography cannot help.

With this in mind, let me ask the reader how you could trust that the above text originated from myself? Of course, “It is printed in the *COOK Report* and Gordon Cook trusts it to be yours!” would be the most probable answer. But should you not be concerned about how Gordon Cook might trust that the above text is my own?

This question is not just rhetorical. Ig-

norning questions like this is at the base of flaws in many Internet protocols, even recent ones. For example, the recently developed SAML (security assertion mark up language) protocol for expressing authorizations in access control and payment systems has no answer to this question. This results in lack of support for audit of assertion dependency between co-operating authorities. For example, suppose that Bob authenticates to the Widget Marketplace using assertion A and receives Assertion B from the Widget Marketplace, whereupon Bob purchases machinery from a parts provider hosted at the Widget Marketplace. The parts provider authorizes the transaction based on Assertion B. If there is a problem with Bob’s purchases at the Widget Marketplace (Bob will not pay his bills) there is nothing in the SAML flow that ties Assertion B to Assertion A. In other words, even though Assertion B has been issued by the Widget Marketplace in response to assertion A, there is no way to represent this information within SAML.

Ignoring fundamental trust issues is not the only problem. One must also emphasize the extraordinary flexibility of the requirements that are served by the use of common names and related identifiers in the Internet. There is a danger here for the creators of protocols that they will (accidentally or deliberately) misuse ordinary, trusted words and familiar concepts in ways that have been artificially restricted by special definitions to fit them for their purpose in the context of formal structures. The result is that they will mislead users as to the actual capabilities of the protocols.

One example is the debate about the word non-repudiation in X.509 and PKIX digital certificates. Even though the only possible meaning of the word is to be something that cannot be repudiated, since X.509 and PKIX cannot provide that meaning but people nonetheless like the “business” sound of it, the word is now oftentimes used to mean merely evidence of authentication.

Misusing trust is also related to the scenario of a spoofing attack. In spoofing, a

user trusts a fraudulent service or information that pretends to be legitimate. Spoofing attacks cannot be prevented by using SSL, digital signatures or encryption. The U.S. National Science Foundation recently sponsored a conference dealing with Internet security issues for voting applications, where spoofing was declared an open and very serious problem without a solution today and long-term research on the subject was recommended.

The problems mentioned above are diverse and touch upon many different aspects of the reliability of Internet communications. We need to solve these problems. A large section of our economy and our lives are already riding on the Internet.

A common question is whether these problems could not be solved by more control. Trust is good but control may be better. However, what to control and where? Unless every user is watched 24x365, or a filtering program is massively used denying functionality to users (as Earthlink does, denying SMTP and NEWS connections to their users, who must then only use Earthlink's SMTP servers notwithstanding the privacy, security and delay considerations), users are pretty much free to do whatever they wish at their connection – including using a different port for a route-around SMTP connection. Savvy users can evade controls. Thus, control does not appear to be effective in an internet. The essential point is that the Internet is a network of networks that has no central control point to be controlled. Not even by means of the DNS.

Here, a different approach imposes itself. Since it is illogical to break communications in order to ensure reliable communications, we ask: can reliable communications depend on trust? And, if so, what is controlled, and how?

Trust vs. Control

In the discussion of trust versus control, it is instructive to view trust as an open-loop control process in control theory terminology, i.e., a control process that

does not rely on a closed feedback loop in order to achieve its purposes. This comparison allows one to recall the advantages and disadvantages of open-loop control vis-a-vis closed-loop control and apply them respectively to trust and close surveillance (also called control).

In control theory, the basic parameter to measure performance is position-error, which translates here to the actual response as compared to its expected or estimated (i.e., trusted) response. In open loop-control, one method frequently used to decrease position-error is to introduce periodic checks of any convenient system variable, not necessarily the control variable. This is not a feedback loop because it is done after the actuation. This method is equivalent to the well-known dictum “trust but verify” – implying the need for a pre-defined policy of checks and balances that can periodically adjust the trust estimator as a function of observed behavior.

Thus, trust can also be explicitly defined as “trust is an open-loop control process of an entity's response on matters of X” or, less precisely but more concisely, as “trust is to rely upon actions at a distance.”

Interesting qualities resulting from this approach to trust in communication systems vs. close surveillance can be exemplified by the just mentioned control theory analogy regarding the main advantages of open-loop control over closed-loop control. These advantages include: simpler systems (hence, better fault-tolerance); immediate response (i.e., nothing needs to be measured in order for it to actuate); easier design (e.g., avoiding probable but unknown pitfalls of complex designs); easier interfacing (i.e., suffers less influence from and also exerts less influence on the rest of the system); modular design (i.e., complete and interchangeable); and less cost.

Where Is The Center? Where Are The Edges?

Using trust tools to solve the Internet problems we see today is, thus, not only

a natural answer in terms of enhancing the very IM2M communication that is failing, but also possibly easier, cheaper, quicker, simpler, more secure and more successful than trying to take control of the Internet.

There is an additional problem that flows from a strategy of control. Strengthening centralized control would make that single handle of control a single point of failure. Strong centralized control also becomes the one basket for all eggs, which everyone wants to possess. These arguments were presented by the author in the April 2000 issue of *The COOK Report* in terms of domain name issues and are also valid here.

An attempt to bring centralized control to the Internet would also need to deal with that vexing question – what to control and where? In the Internet, no single person can tell which networks are included, because no one is there to tell which networks to include and which not to. Any user can add any number of networks to the system. Centralized control is impossible in an internet made out of open-ended networks of networks.

But if the answer does not lie in centralized control, what is the answer to the problems mentioned above? Could Internet control be decentralized? How would this be effective and not generate even more confusion? How would this provide for planetary reliable interpretation of protocols and their messages by machines?

First, one must dispel the notion of “center” and “edges” existing in the Internet. Of course, a network has “edges” and “centers” but the Internet is not a network, it is a network of networks. Surely, one often hears about “edge control” and “center control” and “trust on the edges of the Net” but let's ask – where is the center of the Universe? Where are the edges? Some structures have no edges and no center. The Internet, we must realize, is one such structure. Talking about “edges of the Net” is like talking about the “last page of the Net” – where is it?

The questions above are not even defined. In the Internet, clients and servers are connected with peer-to-peer internet-working capacity, which capacity may or may not be used by an entity. In other words, all clients and servers are able to route IP packets because of their internet-working capacity. So any client or server can become a center of its own network, as well as an end of IP protocol connections. And there is no edge because the Net does not end there, past that edge. At every point, internetworking can be extended to a neighborhood of points all of which lie inside the Net. Further, there is no privileged center that might be called "the center." Thus, there is no meaning that can be assigned to the phrase "the edges of the Net" in the same way that there is no "center of the Net" either.

This also shows that the argument for "trust at the edges of the Net" is a fallacy. An edge that is able to route IP packets becomes a center. And there is no privileged trust location at an edge in the same way that there is no privileged trust location at a center.

Indeed, as we observe the real world and seek to model the trust mechanisms that allow business and human communication to function, what do we see? We do not see a hierarchical trust structure controlling business from a single center. We do not see "edge control" either, and where it was tried it resulted in anarchy. What we see are inter-entity (intersubjective, see Glossary) relationships heavily qualified in many ways. We see internets, networks of networks, a manifold of networks with multiple control boundaries and lacking a common single reference.

Likewise in the Internet, we see a set of edge-edge, center-center (yes, there are many centers), center-edge and edge-center relationships that we can use to induce (communicate) trust. Trust is always formed from relationships between entities because it is induced by, or re-

sults from, a communication event. Such relationships may be unilateral (e.g., I do not know you and I do not know that you trust me) but not singular. There can be no information transfer between a sender and itself, since all is known. The net transfer is zero (see sidebar "Trust Points"). The only possibility for an entity to transmit information to itself is to do so to the future, but then if and only if there has been some loss of memory.

Note that so far we are not yet talking about how we use multiple channels to induce trust, as first we need to establish the need to use multiple channels before explaining how to use them.

Second, we need to direct our attention to the last sentence of the sidebar "Trust Points," quoted from the book "Digital Certificates: Applied Internet Security" by Jalal. Feghhi, Jalil. Feghhi and Peter Williams, where the authors state:

"Models of integrity and trust in a particular space are best left to the communities of interest, governments, and industry forums, which are most familiar with these groups' specific needs."

This sentence reminds us that trust is always local to and is earned in communities of interest. This dispels the idea of a cookie-cutter approach to Internet control, since each community of interest (for example, your company) will have different goals, different control objectives. Increasingly, users want more freedom in controlling their own connections, bandwidth, services and rules of use. Maybe that's why NAT/IPv4 is broadly used and IPv6 is not widespread – even though IPv6 is better in many aspects.

In summary, the answer needed to solve the fundamental problem of Internet communications is trust. Not trust as blind faith, compliance, belief, or ignorance, but trust as qualified reliance on

information through open-loop control. Trust is that which provides meaning to information. Trust is something essentially communicable between machines and humans, something that can flow in our existing TCP/IP, dial-up and other networks.

But there are rules to this communication. Self-assertions cannot induce trust. Client-server communication is not enough to induce trust. We must move to a network model where not two, but four entities need to be in communication in order for trust to be induced.

Why Four?

Trust, as qualified reliance on information, needs multiple, independent channels to be communicated. If we have two entities (e.g., a client and server) talking to one another, we have only one channel of communication. Clearly, we need more than two entities. It seems unreasonable to require a hundred entities.

Editor's Note: This discussion will continue - most likely in the next *COOK Report*.

Glossary

Intersubjective – also called inter-entity; pertaining to more than one entity. For example, making a medical diagnosis is intersubjective because physicians of a same class (i.e., equivalent as observers) diagnosing patients with the same illness (i.e., equivalent as observables for medical purposes) may arrive at different results. The results depend on a patient-physician interaction. A medical diagnosis is thus not objective (i.e., the diagnosis is not the same for equivalent patients and equivalent physicians) and also not subjective (i.e., the diagnosis does not depend only on the physician). The same happens in other cases, most notably in risk assessment.

The Future of Ideas

The Fate of the Commons in a Connected World

by Lawrence Lessig

New York: Random House, 2001, 352 pages

A Review [Summary](#)

The Internet began as government funded experiment in a new kind of communication. For roughly the first 15 years of its existence, it was a run primarily for the US academic research community. It was not commercial. The commercial telecommunications sector did not consider it to be worthy of its attention until the commercialization of the NSFnet backbone in early 1995 kicked off the six year long speculative frenzy that crashed in 2001. During these six years it seemed that the Internet would reshape global communications and indeed the global economy more than the telephone or railroads a century earlier. Turned commercial and subject to tremendous hype as the engine of and means for commerce for the 21st century, the Internet quickly became a strategic corporate target for the telecom, computing and media and content sectors. It was seen as the next great source of wealth and power. As such it had to be exploited and controlled. (For a look at some of the implications of the combination of network and content, see Robert W. McChesney, *Rich Media, Poor Democracy: Communications Politics in Dubious Times*; The New Press, 2000.)

However, the techies who built it quickly pointed out that it could not be controlled. They claimed with confidence that its architecture would not permit control. Lawrence Lessig's great contribution in his first book, *Code and Other Laws of Cyberspace* was to show the geeks that very likely they were wrong in thinking that the Internet could never be controlled. It has turned out that the Internet is a globally a multi-hundred-billion dollar a year enterprise. The Internet became a prize of such magnitude that it compelled huge efforts to figure out how to control it.

As we saw last month in our discussion of the Blumenthal and Clark paper: Re-thinking the Design of the Internet, the end-to-end principles of Internet architecture are under assault from every direction. In his new book, Lessig astutely bases his analysis on this assault on end-to-end as it takes place at what he calls the physical, code and content layers of the network. The end-to-end principles formed a commons over which builders of the pre-commercial Internet agreed they would not fight. Without huge sums of money at stake cooperation and the preservation of end-to-end as a commons was possible. But from 1995 on the Internet expansion was built on a frenzy of jockeying for first mover advantage. As Blumenthal and Clark show the over riding goals of scaling the network and one's participation in it often effectively meant attacking the end-to-end principle with whatever technical devices and fixes were thought to give one an advantage. It was a war for territory and control. As such it meant the commons be damned.

Lessig's Book Laments the Losses

"My central claim throughout [this book] is that there is a benefit to resources held in common and that the Internet is the best evidence of that benefit. As we will see the Internet forms an innovation commons. It forms this commons not just through norms but also through a specific technical architecture. The Net of these norms and this architecture is a space where creativity can flourish. Yet so blind are we to the possible value of a commons that we don't even notice the commons that the Internet is. And, in turn, this blindness leads us to ignore changes to the norms and architecture of the Net that weaken this

commons. There is a tragedy of the commons that we will identify here; it is the tragedy of losing the innovation commons that the Internet is, through the changes that are being rendered on top," [p. 23].

Lessig is right. Destruction of the commons offered by end-to-end also destroys what enabled a succession of creative individuals to build applications like the World Wide Web. The seed stock that enabled the growth of new applications and new industries founded on them is being devoured by the commercializers and their insistence on control. And no one is moving to stop the destruction. Lessig is surprised that there are no good guys. We could be building a commons in spread spectrum. But we aren't. Why? Because there seems to be a "mistaken inference that if some control is good, then more control is better," [p.72].

"The attitude of the most influential in public policy is that free, or common, resources provide little or no benefit. There is for us a cultural blindness-an unwillingness to even account for the role of the commons. As Yale law professor Carol Rose argues, and as I indicated at the start, though 'our legal doctrine has strongly suggested that some kinds of property should not be held exclusively in private hands, but should be open to the public,' we live in a time when the dominant view is that 'the whole world is best managed when divided among private owners.' The very idea that nonexclusive rights might be more efficient than exclusive rights rarely enters the debate. The assumption is control, and public policy is dedicated to maximizing control," [p. 86].

"The story I want to tell is not about the death of innovation generally; it is about

the relocation of innovation from the diverse centralized internet back to the institutions that policed innovation before. The story is about the bureaucratization and capture of the innovation process – relocating it back to where it was as a response to the structures originally enabled by the Internet itself,” [pp. 140-141]. Entrenched interests have been challenged have now struck back. We find this not surprising. Lessig’s work, like that of Clark and Blumenthal, is useful in showing the depth and breadth of the assault to control and regain the old positions of privilege.

Here is his own description of the task he sets: “Changes threaten the power of those now in power; they will work in turn to protect themselves from the changes. In the balance of this book, I want to detail their work to change the Internet, and the legal culture surrounding it, to better protect themselves. Some of these changes are legal; some are technical; and some use the power of the market’But all are driven by the desire to assure that this revolution doesn’t muck things up-for them. There’s nothing immoral in this desire. This is not a battle between good and evil. Stockholders demand that management maximize its income; we shouldn’t expect management to do anything different. But, even if this is “only business” to them, this does not mean it should be “just business” for us. We need not stand by idly as the Internet is changed. They have their interests; we have ours. And for those who believe that the environment of creativity that the Internet produced was worth something, there is reason to resist the changes that I will describe,” [p. 146].

This is quite a challenge. He implies that if we were educated enough and understood what we were losing that we would fight against the losses. Therefore, he sets out to educate us.

“What’s at stake here are two models for organizing a communications network, and the choice for us is which model will prevail. On the one hand, there is the model of the perfectly controlled cable provider -owning and controlling the physical, logical, and content layers of its

network. On the other hand, there is the model of the Internet-which exerts no control over a physical layer beyond the decision to include equipment or not, and which enables the free exchange of content over a code layer that remains open.”

“As the Internet moves from the telephone wires to cable, who should govern? When you buy a book from Amazon.com, you don’t expect AOL to demand a cut. When you run a search at Yahoo!, you don’t expect your MSN network to slow down anti-Microsoft sites. You don’t expect that because the norm of neutrality on the Internet is so strong. Providers provide access to a network that is neutral. That’s the essence of what the Internet means.”

“But the same neutrality does not guide our thinking about cable. If the cable companies prefer some content over others, that’s the natural image of a cable provider. If your provider declines to show certain stations, that’s the sort of freedom we imagine it should have. Discrimination and choice are at the core of what a cable monopoly does; neutrality here seems silly. So which model should govern when the Internet moves to cable? Freedom or control?” [p. 167.] He is a good teacher. His message is filled with reasons to resist. What is missing is a recipe for how we may do this. We found ourselves wondering as we read whether – ever a believer in the power of education — Lessig believes that if his readers just understand they will rebel and some how come together to fight for the benefits of the ‘commons’.

He continues to educate: “The Internet exposes much more copyrighted content to theft than in the world that existed before the Internet. This much of the content holders claim is plainly true. But as I’ve argued, the Internet does two other things as well. First, the Internet makes it possible (if the proper code is deployed) to control the use of copyrighted material much more fully than in the world before the Internet. And second, the Internet opens up a range of technologies for production and distribution that threaten the existing power.”

“In responding to the shock that the Internet presents to copyright law, it is of course important to account for the increased exposure to theft. But the law must also draw a balance to assure that this proper response to an increase risk of theft does not simultaneously erase the important range of access and use rights traditionally protected under copyright law. If the Net creates an initial imbalance, the response by Congress should not create an equal and opposite imbalance, where traditional rights are lost in the name of perfect control by content holders.”

“That was my argument in Code. But now we should add a second concern to that same story: The response by Congress should also not be such as to permit this concentrated industry of today to leverage its control from old world into the new. Artists deserve compensation. But their right to compensation should not translate into the industry’s right to control how innovation in a new industry should develop.”

“Control, however, is precisely Hollywood’s and the recording labels’ objective. In the context of copyright law, the industry has been very clear: Its aim, as RIAA president Hilary Rosen has described it, is to assure that no venture capitalist invests in a start-up that aims to distribute content unless that start-up has the approval of the recording industry. This industry thus demands the right to veto new innovation, and it invokes the law to support its veto right.” [p. 200]

Lessig cites a dissent by California Ninth Circuit Court of Appeals, Judge Alex Kozinski:

“Something very dangerous is going on here. Private property, including intellectual property, is essential to our way of life. It provides an incentive for investment and innovation; it stimulates the flourishing of our culture; it protects the moral entitlements of people to the fruits of their labors. But reducing too much to private property can be bad medicine.” [p. 203] [Lessig] “Why? For the same reasons we’ve been tracking throughout

this book. [Kozinski] Private land ... is far more useful if separated from other private land by public streets, roads and highways. Public parks, utility rights-of-way and sewers reduce the amount of land in private hands, but vastly enhance the value of the property that remains. [Lessig] The state must therefore find a balance, and this balance will be struck between overly strong and overly weak protection.

[Kozinski] Overprotecting intellectual property is as harmful as under protecting it. Creativity is impossible without a rich public domain. [Lessig] But is that unfair? Is it unfair that someone gets to profit off the ideas of someone else? Says Kozinski, No. [Kozinski] "Intellectual property law assures authors the right to their original expression, but encourages others to build freely on the ideas that underlie it. This result is neither unfair nor unfortunate: It is the means by which intellectual property law advances the progress of science and art. We give authors certain exclusive rights, but in exchange we get a richer public domain."

[Lessig] This balance reflects something important about this kind of creativity: that it is always building on something else. [Kozinski] Nothing today, likely nothing since we tamed fire, is genuinely new: Culture, like science and technology, grows by accretion, each new creator building on the works of those who came before. Over protection stifles the very creative forces it's supposed to nurture. [Lessig] This balance is necessary, Kozinski insists, "to maintain a free environment in which creative genius can flourish." Not because "flourish [ing]" innovation is the darling of the Left, but because innovation and creativity were the ideals of our founding republic," [204].

Lessig is one of the very few commentators who understands the significance of spread spectrum wireless as a tool for enabling grassroots communications. Through talking to folk like Dave Hughes, he has also come to understand the public policy shortcomings of the federal government's embrace in the 1990s of

sale for billions of dollars of private property rights to large corporate customers. It gives the large licensees huge reason to resist the use of spread spectrum as an unlicensed commons.

"As Dave Hughes has asked: [W]hy should AT&T, who is offering a wireless service.... consent to competitors in their own area? ... [I]t's not just a question of interference now. Now it becomes ... opening the door by their consent to competition. And the last damn thing big companies want is competition."

"The danger in selling spectrum or, more precisely, in not experimenting broadly with unlicensed spectrum is that existing spectrum users will be able to use purchased spectrum to resist changes in spectrum policy that might threaten their business models. By selling spectrum now, before alternative uses can be developed, we create a world where the resources for these new alternatives are held by those with the strongest incentive to stop them. As Eli Noam puts it, it is like "having the old AT&T auction off the right to compete against it. Under such a system, MCI would not have emerged."

The concern is not just about spectrum owners; it is also about the nature of the existing spectrum uses. The dominant and fastest-growing spectrum use right now is mobile telephone systems. These systems are architected in just the way the old telephone network was- intelligence is located not at the ends, but instead in the network itself The cellular phone companies retain control over how the cellular technology develops; if you want a new application for your (increasingly powerful) phone, you will get it only if the telephone company wants you to.

"This architecture for a wireless system creates the obvious protectionist risks. And as Charmed Technologies CEO Alex Lightman puts it, we are already seeing these risks mature into protectionist practices." "[T]here is a nice little cozy menage a trois between the companies that are providing infrastructure and the companies that provide the

handsets, and the monopoly carriers or the oligopoly carriers."

"By selling the spectrum, the carriers have a strong incentive to assure returns sufficient to recover the investment in spectrum. These returns are best assured (or at least it seems to the companies that are best assured) if the companies husband the market power that is carried over from the noncompetitive telephone world (recall: much of the action here is international, where competitive phone systems don't yet exist). Thus, the willingness of the existing players to open up their spectrum to a wildly different form of use-one that would be much more competitive-is unlikely at best. Recall the words of AT&T executive Jack Osterman in 1964: "[W]e'll be damned if we allow the creation of a competitor to ourselves."

"A policy from the FCC that does not create a strong opportunity for an alternative to develop is designed to protect existing interests. To not encourage or permit wide-scale experimentation, to not set aside much broader unlicensed spectrum, to protect existing uses against any interference these are policies designed to preserve the old against the new. They are just what we would expect from government regulation of spectrum; they are much less than we should demand after the experience of the Internet. The government's role should be to induce investment where there is a great deal of social value to be created. This is precisely the opportunity with unlicensed spectrum,"[pp. 226-27]. It seems to us that the goals of CaNet*4 also fall into this category.

Oh Where are the Commons Protectors?

Let's summarize. Lessig recounts some parts of our legal and economic system that still preserve things like highways as a commons requiring no prior permission to drive on. He adds that such is the case because there is a general recognition that having it so increases the overall wealth of society. He raises an eloquent argument as to what SHOULD be done. He complains that we are destroy-

ing long term gains in order to reap short term benefits. He provides an excellent summary of the most egregious of recent court cases where owners of old content have successfully used the courts to make sure that their property rights are extended to all aspects of the Internet and far beyond what we might have deemed possible only five years ago.

He suggests that the government could play a role in providing infrastructure and separating content from distribution. But he has no clue where such “good behavior” might come from. We need it. If only we had it . . . He talks about an American cultural blindness in allowing the scope of what is defined as intellectual property to get out of control. Large corporate interests are changing the architecture of the Internet in order to control it, he laments. He points out that they are acting to protect old entrenched backward interests. No one is coming forward to stop them he complains. We are losing future innovation. We should wake up and demand that it STOP!

He says what regulators here should do. “[R]egulators should begin to evaluate changes to the network in terms of the neutrality of end-to-end.” True. But he also seems almost mystified that they have not done this. He states what regulators and other policy makers have just done in Canada. “The Canadians have required open access for broadband providers and a recent Report of the National Broadband Task Force has endorsed a “bill of rights” for broadband users that assures continued consumer choice for any build-out of a network employing government funds. This choice of policy has apparently not harmed Canadian access. According to a recent OECD report, broadband connections are twice as common in Canada per capita as in the United States,” [p. 249]. We conclude that the Canadians have not yet sacrificed all ideas of Canadian national and public interest on the altar of the unfettered free market. Won’t it be interesting if what is derisively referred to as “socialist leanings” in many US political circles creates an environment in Canada that is a fertile seedbed for new technology development?

We think his arguments are correct but we found them intensely frustrating. Can he really be so naive as to believe that it is just technology blindness that is preventing the right choices from being made? What Lessig does persuasively document is the myriad of ways in which the apparent decision forcing power of the bottom line has channeled corporate decision makers to pursue policies of control that are designed to squeeze every possible cent from corporate owned assets. He ends his book on a note of considerable pessimism.

“Additions that benefit either company [Microsoft or AOL/Time Warner] will be encouraged; additions that don’t, won’t. We will have re-created the network of old AT&T, but now on the platform of the Internet. Content and access will once again be controlled; the innovation commons will have been carved up and sold.”

“This is the future of ideas. It could be different, but my sense is that it won’t. If we were more like Hatch, more skeptical of “gatekeepers,” whether private or public; if we were less like Jay Walker, eager to view every government-granted right as a God-given property right; if we were more like Richard Stallman, committed to a principle of freedom in knowledge and to a practice that assures that the power to control is minimized; if there weren’t so few Paul Barans, willing to struggle for many years to force a monopoly to face itself- if all this were so, there would be reason for hope.”

“But we are not. We are a democracy increasingly ruled by judges. We elect a Congress that is increasingly chained by lobbyists. And we are a culture that deep down believes in this counterrevolution: that strangely thinks that this increase in control makes sense.”

“As commentator Gordon Cook writes: The Internet revolution has come and gone. It has created a tremendous burst of innovation [-a] burst that now looks to have been mismanaged.... [T]he people who did the least to advance the new technologies seem most likely to control

them. We are left not with the edge-controlled intelligence of the [end-to-end] network but with the central authoritarian control of the likes of AOL Time Warner.” [He cites our “What Good is the Stupid Network. . .” article from the September 2001 issue.]

“The irony astounds. We win the political struggle against state control so as to re-entrench control in the name of the market. We fight battles in the name of free speech, only to have those tools turned over to the arsenal of those who would control speech. We defend the ideal of property and then forget its limits, and extend its reach to a space none of our Founders would ever have imagined.”

“We move through this moment of an architecture of innovation to, once again, embrace an architecture of control -without noticing, without resistance, without so much as a question. Those threatened by this technology of freedom have learned how to turn the technology off. The switch is now being thrown. And we do nothing about it,” [pp. 267-68]

Final Thoughts – Why the Commons Died

We found ourselves staring at and pondering the sentences that explain why he is not hopeful: “We are a democracy increasingly ruled by judges. We elect a Congress that is increasingly chained by lobbyists. And we are a culture that deep down believes in this counterrevolution: that strangely thinks that this increase in control makes sense.”

This last sentence about belief in counter revolution and in the desirability of increased control puzzled us. We thought long and hard and found ourselves writing the following paragraphs. They are deeply personal and in that sense we are not sure that they belong here. We have asked ourselves in the aftermath of September 11 where we are going. We think that others are doing the same. The answer is obviously deeply subjective and personal. Perhaps it is time to take a big step back and ask questions for which we earlier may have felt no need to ask? We don’t have a prescription for change but

without even an awareness that we have a problem certainly things will not change. Lessig asserts we have a belief in counter revolution. If he is mistaken, then what do we have? Here is our take.

In universities and government funded labs in the 60s and 70s we created packet switched computer networks and built the Internet on a platform of edge empowerment. It carried with it the potential to give end users capabilities impossible in the centralized world of the mainframe. But, as this technology matured, in a series of parallel but otherwise unrelated developments, we allowed our politicians to trash the idea that there could any longer be a public sector worthy of allegiance and respect. Government became seen as a necessary evil that really had to exist only for the national defense. Meanwhile the only goals that were truly worthy were those that freed the 'market place' to be the compass and guiding light for our society. In the process, what we have lost I would argue is much more than the capacity to see the wisdom of maintaining public commons and the ability to understand that placing limits on private corporate reach may preserve the ability of the internet to function as a platform that supports further technology innovation.

We have lost the ability to distinguish between public and private good in our political life. We have lost this because of the political success over the last 40 years of the idea that the free market is the most important goal of life. Our country was founded 225 years ago on the idea that government should serve the individual. And that if governments failed to do this, the citizens should be free to rebel and establish a new polity that did render possible for each citizen life, liberty and the pursuit of happiness. The French followed through with their own revolution a dozen years after ours.

At the beginning of the reign of Louis XVI public life in France was focused on the idea that the great landed aristocracy should work in tandem with the French throne to preserve "stability" by declaring the thought of the enlightenment to be seditious. The wisdom was that the

aristocracy was the rock of society from which wealth flowed. Intellectuals of the day had no right to declare the sanctity of the individual. During the 1780s the influence of our own revolution became the ideological underpinning for the mounting by the growing French commercial class of a successful challenge to efforts of the king and aristocracy to maintain its privilege.

Our current problem may be that we have reverted to an environment - perhaps unwittingly - where our lives have been usurped by a new "King George," and by a new aristocracy. That aristocracy was a new group of conservative politicians. Pitting its own ideological worship of the free market against a subversive threat and enemy called "communism," it successfully swept our political system in a direction that glorified the power of the corporation and its ability to serve as an engine for the creation of wealth. Undoubtedly, well meaning at the time, we contend that, flush with their own success, these folk are more and more focused on their own self-preservation. [See for example <http://www.heritage.org/shorts/20011101stimulusletter.html>]

The pendulum has swung to such an extreme that the captains of industry, as they create their new structure of global capitalism, have created an environment where they define the national social political and economic agenda in terms of what benefits their bottom line. Our government no longer rests on its early foundation based on enlightenment philosophy and a goal of the empowerment of individual. Corporations rule our political and legal system. With no appreciable influence left in that political system, we live as new serfs under a government enforced legal and economic system where corporate owned politicians call the shots and let us pretend that elections still matter and can change things for the better. We have been overthrown and the quality of public discourse is so low and so fractured that we don't even know it.

Add to all this a failure of our public education to define and instill any deeply held faith about what might constitute the

public good of this nation. Season the result the triumph of television - advertising fueled psychology that far too often the higher purpose in our lives is to acquire more and more of the physical goods that our factories produce. Can this be what Lessig means by a cultural bias for counter-revolution and control?

We sent an earlier draft of the above paragraphs to Larry Lessig saying we were puzzled by his conclusions. He replied on November 8, "Gordon, I don't understand the puzzle. As a careful and complete reader of me, you know my view about corruption in the political system. Both in this book and the last, that was an underlying complaint. The difference in this book is that I need to explain the "corruption" of the judges too - which is not money corruption but something else. I tie it to something about our naive attitude about property. You say I struggle not to be labeled anti-property: but that's because I am not. It is not a pose, or strategy: it is truth. My argument is not the standard "anti-property, pro-commons argument"; my argument is a pro-property, pro-commons argument-but just a conception of property real to our tradition. I don't think the reason we are here is just because lobbyists have bought Washington-that's an important reason, but it doesn't explain enough. I think we are here because most people think the argument the lobbyists make (maximal IP, etc.) is a good one for society as a whole. Not just Disney, but everyone. That is a harder corruption to understand than the standard special interest complaint. But in my view it is a more important corruption to understand as well."

There Can Be Some Cause for Hope

To the extent where they have acquiesced the sole criteria profit maximization, corporations have failed us. But we all have to acknowledge that there are many indirect ways that responsible citizens and corporations can benefit a larger society. For example, corporations can foster the public good of a healthy and well-educated population, which respects a broad range of property rights. A society on a

permanent war-time footing may display an instability that actually hurts the corporate bottom line. Perhaps therefore CEOs could fulfill their duties to stockholders by exploring an approach that including the use of property in order to promote a general public good that comes with the understanding that along with property rights and responsibilities must be bundled. In short, a body politic concerned with increasing viability local self-sustaining communities that are the foundation of political and economic stability.

Postscript: A Pointer to Mid November Lessig Paper at Duke Law School

From November 9 to 11 Duke University Law School held a remarkable Conference on the Public Domain
<http://www.law.duke.edu/pd/papers.html>

The papers are well worth looking at.

Lessig called his talk; The Architecture of Innovation
<http://www.law.duke.edu/pd/papers/lessig.pdf> It is a powerful summary of the major arguments of his book.

“We have this view—this taken for granted, background view—because for the last hundred years, we’ve debated a related question, and that debate has come to an end. For the last hundred years, the question exciting political philosophy has been which system of control works best. Should resources be controlled by the state, or controlled by the market. And this question, we all rightly believe, has been answered. In all but a few case, for a wide range of reasons, we know this: that the market is a better tool for controlling resources than the state. That between the two, there is no real debate. The communists roll on the dustbin of history. But this confidence obscures a distinct and more basic question. This certainty about the market over the state leads us to ignore an issue that comes before. Not the question of which system of control is best for any given resource; but

instead the question – should a resource be subject to control at all? Not the market vs. the state, but controlled vs. free. If communism vs. capitalism was the struggle of the 20 th century, then control vs. freedom will be the debate of the 21 st century. If our question then was how best to control, our question now will become whether to control. What would a free resource give us that controlled resources don’t? What is the value in avoiding systems of control?”[p. 178]

Lessig’s questions about control are valuable. Appropriate for war and geopolitics they are also appropriate for the Internet. The personal computer began the decentralization of control from the machine room to the hands of end users 20 years ago. We have seen earlier in this issue how CANARIE intends to make personal bandwidth possible by leveraging the commodification of lightwaves from long term centralized control by carriers to short term peering made possible by giving users the ability to set up and tear down their own networks at the physical level.

Lessig at the end hammers a subtle but important point. It is one that harkens back to the analysis that we made in our comments at the end of our review when we lamented the capture of our politics by those who hold the ideas of free market and property rigidly sacred. Thus if some property is good more must be better. Life is turned into a race to acquire and the corruption of the lawyers that Lessig rails against has come apparently from their being willing to “propertize” copyright turning what was originally a government backed monopoly into intellectual property which according to our ruling ideology is sacred and hence something to be extended, conserved and controlled.

“Ideas that are taken for granted; that are unquestioned in this culture; that to question, would render you an alien; these ideas are the ideas that will make control the future. For these ideas take for granted the property in intellectual property; these ideas have lost the distinction that our framers made clear—by speaking as they did, not of intellectual property, but of monopolies and exclusive rights.

That’s what a copyright or patent is—a government backed monopoly, not over a rivalrous or scarce resource like land or apples or heated homes, but over a nonrivalrous resource that the enlightenment taught us should be shared among more than the church. IP is not P, but this truth is lost on us.”

“And so deeply is it lost that we don’t even notice the irony it produces. We speak of a commons as if it is only a tragedy; we recall the public domain as if it were simply an echo from some romantic past; we embrace, as Professor Rose says, the idea that the whole world is best managed when divided among private owners, and we proceed to divide the world among private owners. Most Americans agree with the Disney Corporation that Mickey Mouse is Disney’s now and forever; they don’t even notice the irony then when Disney can make millions off of Hugo’s creation, the Hunch Back of Notre Dame, or Prokofiev or Pocahontas. So invisible is public domain that we don’t even see it when it is everywhere around; so invisible is the idea that the free might matter to creativity, that when it is enclosed, we are convinced this is progress.”

“Our future is this: the free speech clause of the first amendment will be read to entitle those who own the wires to change the logical layer and make it owned as well; the free competition principle of the Sherman Act will be read (by the same circuit we might notice) to entitle the owner of the platform that most affects this logical layer (that one company whose name I have not uttered) to code that platform to discriminate as it wants; and the free culture that we have seen flourish in this commons built by the Internet will be captured and controlled again by those who control most of the content, and by those who succeed in congress in expanding their control from the imperfect to the perfect.”

“The future of control will get built by an idea; the idea that property is good so more property is better. It will get sanctioned by a culture that has forgotten any distinction, and that is so blinded by what it has forgotten that it does not even notice when the most extraordinary innova-

tion that our culture has seen since Thoreau was a name most Americans could spell is built not on an architecture of perfect freedom; not in a world where every layer is in the commons; but also not on an architecture where control was the rule; not on an architecture where every layer was owned: but instead on an architecture that mixed freedom and control; that built property within a commons; that got its life from this mix of property and the commons." [p. 189]

"We allow these changes, they don't just happen. We stand back as they occur, they don't happen in the night. We let them occur because most of us believe they should; control is good, better control is better, these systems of control are ways to make sure the better comes from

the good. It is an attitude and blindness and a pathetic resignation that permits this change. So enamored we are with the invisible hand, so convinced we are of the genius of property, so blind we are to what makes innovation possible, that we allow the undoing of the most significant chance for something different that we have ever seen." [p. 190]

"These are people who have not been to Duke. And so I come to Duke to do little more than report on a war we are losing. Of a culture that can't see the potential that this architecture presents. Of a politics that scorns anyone who questions that uber vision of perfect control. The irony astounds. We win the cold war against state control so as to reentrench this system of control in the name of the

market. We fight battle in the name of free speech, only to have those tools turned over to the arsenal of those who would control speech. We defend the ideal of property, and then confuse its limits, and extend its reach to a space none of our founders would ever have imagined. We move though this moment of an architecture of innovation, to once again an architecture of control. Without noticing; without resistance; without a question. This you may notice is a contradiction in our tradition. You might be tempted to then repeat my favorite line from Jamie's book, "I have nothing against contradictions, some of my best friends are contradictions." This is a contradiction we should be against. Yet, we, Americans, are not." [p. 191]

ICANN Annual Meeting Security Sideshow Fails to Upstage Completely the Nasty Question of an At-Large Membership

Editor's Note: ICB Toll Free News is an online publication that takes a critical look at 800, DNS and ENUM from behind-the-scenes. For more information see <http://JudithOppenheimer.com>.

SEX, LIES, AND FLUID DEFINITIONS

Marina del Rey, CA (ICB TOLLFREE NEWS) On November 8th, ICANN President/CEO Stuart Lynn announced during a press teleconference, "No decision could be taken on the ALSC [Report at ICANN's November 15th Board meeting]. ICANN's ByLaws require a period for comment from it's Supporting Organizations and that would be the case whatever the agenda next week."

One week later, on November 15th at the ICANN Board meeting in Marina del Rey, ICANN Chairman Vint Cerf <<http://cyber.law.harvard.edu/scripts/rammaker.asp?s=cyber&dir=icann&file=icann-111501b&start=3-19-30>>stated otherwise, apparently as unconcerned with the ByLaws as the ALSC itself: "We do have intentions to move on the Report in this afternoon's meeting," he

confirmed replying publicly to ex-ICANN Chair/current At Large Study Committee member Esther Dyson's smarmy ploy for immediate Board approval of the ALSC Report.

Further into the public meeting, Dr. Cerf was good enough to read aloud my <<http://cyber.law.harvard.edu/icann/mdr2001/archive/subcomments-111501.html>>remotely posted question seeking clarification of this obvious contradiction, "whether referring to the "basics" or the details of the Report," between Lynn's statement, and his own.

"I meant to say that the subject would come up in the afternoon Board meeting," Cerf <<http://cyber.law.harvard.edu/scripts/rammaker.asp?s=cyber&dir=icann&file=icann-111501b&start=3-19-30>>adlibbed, "not that we would make a final decision on the ALSC Report."

Yet what "came up" at the afternoon Board meeting was not "the subject" but an already prepared formal ALSC Resolution, "that the President and CEO is directed, in consultation with the ALSC, to begin ... planning for carrying out an At

Large election process in 2002 consistent with the recommendations contained in the ALSC Final Report..." .

Now, ICANN can split hairs all it wants over the meaning of "final decision," but the artifice is all too reminiscent of President Clinton's fluid definition of "sex."

It's all wet. And you get screwed.

Background

When Commerce first sought to transition Domain Name System management to a publicly accountable "bottoms up" private organization, ICANN pledged to create an open membership structure to assure public oversight and public input. ICANN promised to create an "At Large" membership which would directly elect 9 members of the Board, as a counterweight to the 9 directors elected by the industry-based Supporting Organizations. See

<<http://www.ntia.doc.gov/ntiahome/press/ICANN111098.htm>>Letter of Esther Dyson, Interim Chair, to J. Beckwith Burr November 6, 1998. That letter contained the following promise from the ICANN Board:

"Some remain concerned that the Initial Board could simply amend the bylaws and remove the membership provisions that we have just described above. We commit that this will not happen. In addition to our commitment, the U.S. government has publicly stated that it will maintain oversight during the transition period, and we fully expect that the creation of a membership and the transfer of authority to a fully elected Board will occur before that transition period ends." (emphasis added)

On the basis of this among other obligations, the Department of Commerce entered into a cooperative agreement with ICANN.

ICANN repeatedly promised that establishment of an open membership, direct elections, and the resignations of the initial Board members were it's "top priority." See, e.g., <<http://www.icann.org/correspondence/icann-to-doc-19july99.htm>>Letter of Esther Dyson to J. Beckwith Burr, July 19, 1999. In particular, Esther Dyson again assured the Department of Commerce (and the Internet community as a whole) on behalf of ICANN that:

"Our goal, which I know you share, is to replace each and every one of the current Board members as soon as possible." (emphasis added)

ICANN made similar pledges to Congress. In sworn testimony before the Congressional oversight hearing on July 22, 1999, Dyson testified:

"As to the second wave, it is ICANN's highest priority to complete the work necessary to implement a workable At-Large membership structure and to conduct elections for the nine At-Large Directors that must be chosen by the membership. ICANN has been working diligently to accomplish this objective as soon as possible. The Initial Board has received a comprehensive set of recommendations from ICANN's Membership Advisory Committee, and expects to begin the implementation process at its August meeting in Santiago. ICANN's goal is to replace each and every one of the current Initial Board members as

soon as possible.(emphasis added)

<<http://www.icann.org/dyson-testimony-22july99.htm>>Testimony of Esther Dyson, Chair, ICANN, before the House Commerce Committee, Subcommittee on Oversight and Investigations, July 22, 1999.

Then just one month later, in August 1999, the initial Board members extended their terms another year, and adopted a resolution to prohibit direct elections of directors by the At Large membership. See <<http://www.icann.org/santiago/santiago-resolutions.htm>>Santiago ICANN Board Resolutions.

In March 2000, ICANN finally created a mechanism for a general membership, but it would only allow the election of five of the promised nine At Large directors. See <<http://www.icann.org/minutes/prelim-report-10mar00.htm>>ICANN Board Resolutions. In addition, the Board again extended the terms of initial Board members, allowing four of them to remain on the Board until October 2001.

In August 2000, the Board adopted a resolution calling for a "clean sheet" study of the At Large membership, bringing us to today. Among the members of ICANN's At Large Study Committee, the now ex-chair of ICANN, Esther Dyson.

The subsequent At Large Study Committee Final Report recommends six At Large members and 12 industry members, effectively neutering public participation and any hint of democracy in Internet governance. ICANNWatch dissects the At Large Study Committee Final Report <<http://www.icannwatch.org/article.php?sid=458&mode=&order=0>>here, noting among other things, that it relies on argument by assertion; it is arbitrary, dishonest, naive, vague and exclusionary.

It is also <<http://www.icannwatch.org/article.php?sid=460&mode=&order=0>>widely rejected worldwide.

In contrast, an independent NGO and Academic ICANN Study (NAIS) was simultaneously conducted by experienced

researchers from nine organizations worldwide with substantial expertise in ICANN. (See

<<http://www.icbtollfree.com/article.cfm?articleId=5429>>IF YOU CAN COUNT, ITS A NO-BRAINER.) It was well received by a broad range of Internet constituencies. And ignored by ICANN.

Copyright © 2001 ICB, Inc.

Editor's Postscript: On November 20, in BWG: **Attorney:** Perception counts for a lot, and if the perception that the ALSC is headed in "the right direction" gains any traction, good luck trying to steer it off its current path.

Ted Byfield: the ALSC is headed in exactly and only the direction it wants, and I'm very skeptical that perceptions one way or another will change that course. hence the need for spin: it's not just recovered pride, it's preparation to forge ahead as if nothing happened. Do you think ICANN's spin artists will convince the board that they didn't have the conversation they had? or convince the civil society crowd that the ALSC report is good? As to the latter, we've seen what happened with the Cairo compromise, and that is exactly what will happen--indeed *is now happening*--in this case.

And as long as the civil society crowd continues to try to develop an ALM *within* the ICANN process, they'll get Cairoed, pure and simple. Staff will interpose themselves as gatekeepers and produce, variously, inaction and simulacra, and the ALM will end up a toothless and marginal ghost, if it ends up at all. That's why the ALM *must* be a constructed outside that process, as a separate legal entity.

Different Attorney: They'll never, EVER let that happen. Best to let them do what they're going to do (i.e. castrate the ALM) and then challenge their 501(c)(3) status because 1) it needs to be re evaluated because of a material change to what they promised they would be and 2) it has now become a trade organization since users don't have equal footing which was one of the reasons given for being tax exempt. California, and or the IRS will be quite interested I'm certain.

Byfield: I have a slightly different scenario in mind, but we're heading in the same direction, yeah.

Interview / Article Highlights

Ca*net 4, pp. 1 - 18

[Go to Executive Summary](#)
[Go to Full Interview](#)

St. Arnaud: To light your fiber you have to make a very substantial investment in equipment. And therefore telcos become very interested in finding new applications and services that can drive potential customer demand.

That problem started us thinking about customer owned networks, object oriented networking and OBG. Today networking is like computing was 40 years ago when the market was dominated by large mainframe computers. But in the 1970s the mini-computer came along followed by the PC which fundamentally changed our thinking of how to do computing. Computing became personal. The user was empowered to develop new applications and services that were not possible on a mainframe computer. With CA*net 4 we hope to move networking in the same direction as computing has gone in the last 30 years.

However, I stress again that OBG is experimental. For some unforeseen reason it may not work or have some other deficiencies that preclude its deployment. This will not prevent us from still pursuing our vision of customer owned and controlled wavelengths. [p. 3]

St. Arnaud: From day one we will be assigning ownership and control of individual wavelengths or STS channels to the GigaPOPs, universities and perhaps even individual researchers. They will be free to trade and swap amongst themselves and do what ever they want with those wavelengths. From day one we will also encourage these organizations to directly peer with each other and other international research networks using these wavelengths. But, initially the BGP optical peering will be done manually. Once OBG is successfully implemented, it

will allow these organizations to automatically change the routing of the wavelengths and peering relationships without first contacting CANARIE. So rather than operating a traditional hierarchical IP network as many other research networks do today, CANARIE will only offer an aggregate IP network as an optional service for those organizations that don't need their own wavelengths. [p. 4]

St. Arnaud: Increasingly in large IP networks, routing is done at the edge with switches placed at the core. This is not a new concept. Everyone is doing this, even the telcos. But, a question on which views diverge is who controls the switch in the core? Right now, the telcos and most of the telco suppliers are saying that the switched core should be owned and controlled by the telephone company. The customer can have routers at the edge, but the carrier will go back to its roots and do what it knows best, which is switching circuits across the core. In this case it will do it by setting up an MPLS path for the customer and, of course, charging him accordingly.

We are saying let's try something different and let the customer own the wavelength across the network. This is the big difference. This is why we state that switch architecture for CA*net 4 must allow external users to manage cross connects, provision user VPNs across the switch, and so on. [p. 5]

St. Arnaud: We co authored a draft with Viagenie and took the result to the IETF. One of the challenges we ran into right away was that the optical working group in the IETF said: to get this to move forward, you need an endorsement by a major carrier. We said a major carrier is not going to endorse this because they perceive it to be against their own interests.

We will not go away and ignore the IETF, but the problem is that to implement OBG within a router we need to get

Cisco or Juniper to endorse it and include it within their code. To do all that they first have to see a business case and so on. It could be ten years under those conditions before we see an implementation and an IETF standard. So when Wade pointed out the Bill Joy article and said lets use agents as network objects, I was interested. If for we do it this way, we can implement it and interface with existing BGP processes. With an agent approach we can implement this by using open standards. [7]

St. Arnaud: One of the problems that we have been struggling with is that there is a real telco mind set among all the vendors. It is very hard to get them to understand some of the principles we are trying to develop here. Consider the traditional Network Management System. It assumes that each carrier is going to manage his network cloud into which no one else can see. One of the popular trends is for something called CMN or Customer Network Management views. They have a big management system for themselves and they may give you via the web a little peek into the portion of it that belongs to you. But you cannot yourself directly do anything with it.

The vendors keep coming back to us and asking if we want this customer view capability. We say no. That is not what we want. A customer window must actually directly enable customers to change the network that it shows. While we think that our concept of OBG and customer owned wavelengths seems straight-forward and simple, we keep being surprised by how difficult it seems for the vendors to be able to grasp it.

We have been saying that the management system needs to be thought of as having four layers. First, there is to be a physical partition layer. One where you partition the switch. The analogy I like to give is that this should function like a carrier hotel where each customer gets its own room in which it does what it wants.

[pp. 9 -10]

St. Arnaud: ..these cross connects can be incorporated into the user's routing domain. By analogy in the Internet today you have dozens of exchange points in North America and dozens more scattered around the world. If you are an Internet service provider, you have many connections to different Internet exchange points. What we are attempting to do is to extend the concept of these interconnect points or IXs down another level so that they are accessible by individual customers.

Rather than having a few hundred IXs for a small number of big ISPs, we will strive toward the point of having thousands of IXs for individual institutions and users. Ultimately, as an end user, rather than depending on your ISP to interconnect you to somebody else, the intent is for you to be able to cross connect to that person just like the two of you were actually interconnecting at an exchange point. [p. 11]

St. Arnaud: We expect the winning vendor to partner with CANARIE in developing IRR tools and NMS on the switch. We hope that within six months we will have something to experiment with a on the switch and that in a year or so we may have something that works really well. But let me emphasize again that this is research. It is conceivable that at the end of the day we may conclude that it was a dumb idea. But we hope certainly that ultimately we will say: this is a really neat product and that the switch vendor will then take it out into the commercial world. [p. 12]

St. Arnaud: For example, let's say the regional network in British Columbia (BCnet) is given control over four pair of national wavelengths from CANARIE. BCnet then transfers control over of two pair of the wavelengths to the University of British Columbia which in turn gives one pair to TRIUMF, the high energy physics research facility at the University. TRIUMF now has its own national network cloud with a pair of wavelengths that can run on each CA*net 4 switch. Initially, TRIUMF may connect to SNO, another physics research facility, at Sud-

bury. A bit later it could decide to switch its wavelengths to the STAR LIGHT to connect to Argonne Labs in Chicago. Or it could chose to drop off a portion of its wavelength as an STS channel at Calgary and another in Sudbury. All the time the topology and routing of the wavelength is under control of the researchers at TRIUMF. The decentralization we envision will give our customers extraordinary flexibility in what they do with their resources. [p. 14]

COOK Report: It seems that the Internet is continually in search of a business model. Perhaps by 1996 or 1997 we found a viable business model for dial up and leased lines. But by 1998 DWDM and optical gear were making their impact felt. Over the next two years they threw all earlier assumptions out of whack. Could it be that the Internet and telecommunications in general now needs a new business model that takes the optical revolution into account? While the problem of the local loop still looks especially nasty, Ethernet in the first mile development looks very interesting. Also there are new ideas emerging about how to get fiber to the home. See for example: http://lw.pennnet.com/Articles/Article_Display.cfm?Section=Articles&Subsection=Display&ARTICLE_ID=0252

What are your thoughts on what this model might look like and how it might work itself out over the next five years?

St. Arnaud: I think the business and architecture model of the future will be of control and management moving increasingly closer to the edge, not only of the in terms of applications, but also in terms of control over the infrastructure. I think that one of the drivers for this will be as a consequence of the issues that of Larry Lessig has raised where content and distribution companies are trying to exert control over the Internet infrastructure to protect their intellectual property interests. Decentralization and minimizing control at the center will help thwart these challenges. We are already working on concepts with our industry and research partners to extend this concept of customer control of wavelengths all the way to the individual home.

For example, one of the concepts we are exploring for delivering broadband to the home is where the homeowner owns and controls individual strands of fiber. We are investigating how to extend the principles of OBG and distributed computing to the consumer level through concepts such as Reverse Passive Optical Networks (RPONs). With RPONs the control of the wavelength and fiber routing, including moves, adds changes is always under the control of the end user. From time to time the end user may delegate that control to the service provider of their choice, but the fundamental principle remains the same - end user ownership and control at the edge.

In the future we see a physical network infrastructure that closely parallels Morpheus and other peer to peer networking paradigms. The end user will have a choice of whether they wish to subscribe any number of "walled garden" service providers. Or they may chose to physically connect to community networks to share files and data with high speed Gigabit wavelengths bypassing all traditional hierarchical service providers.

A lot of these concepts are still very speculative and unproven at this point in time. But what it does point out - is the critical role that research and education networks still play in the ongoing development of the Internet. The Internet would have probably never been developed if we restricted ourselves to the telecommunication business models of the mid 60s and 70s. The Internet only came about because universities and research centers began to explore new ways of interconnecting computers. In the late 1990s the conventional wisdom was that university research networks had done their job - the Internet innovation phase was over and it was time to focus on commercialization and scaling issues. But in reality the Internet revolution is far from over. Most importantly research and education networks will continue to play an absolutely critical role in exploring new concepts in networking which initially may appear very radical to the traditional telecom world.

[pp. 17-18]

Trust as Qualified . . . , pp. 19 -25

[Go to Executive Summary](#)
[Go to Full Interview](#)

Editor's Introduction

Trust is a word that is commonly applied to many situations and consequently has many shades of meaning. The following essay by Ed Gerck focuses on one precise set of coherent meanings: the concept of trust in the context of communication. More specifically, in the context of the engineering problem of Internet communications. At the same time he demonstrates why trust is needed in this context. Trust is considered something essentially communicable, but with specific rules for its communication. Gerck's exposition also discusses the induction (communication) of trust in heterogeneous environments, from human to machine, machine to machine, and machine to human.

By allowing trust to bridge the many gaps between human and machine, people will be able to tailor their own human to human communication needs via the Internet. What this means is that communities of interest, as networks of people, can build their own networks within the Internet according to their needs, without any limitation imposed on them by their Internet connectivity. [snip]

Steffrud: In short, the Internet does not really have a center or edges. It only has connection points, each of which can be connected to any other such connection point for the purpose of packet exchange. One reason that the Internet does not have an edge (as I just realized) is that at any termination connector, it is possible to extend the Net beyond that point by relaying packets, or by relaying messages, via dial-up modem, FAX or printer, or word of mouth, for that matter. So we suddenly discover that we cannot define any edge of the Internet.

Gerck's communication concept of trust may be just the beginning of a broad-based understanding of a new view of the

Internet where security is a core part of the design. This new view of the Internet is that of a large collection of local network centers and edges, of potentially overlapping subsets of the total Internet. It is built around local common interests and purposes (communities of interest), but with a global communication pattern that closely resembles how we humans communicate across such boundaries and how our commerce works; and it looks like an assembled collection of networks, each of which has its own local centers and edges when we observe them closely, but the global collection of these local networks into the whole Internet doesn't have a single global center or any edges.

Now, since we have so many available connections, Gerck is saying, let's use sets of connections to enable us to transfer trust using distinctly separate multiple channels. Except that, in the essay that follows, he leaves for a next article the discussion about how one can use those multiple channels to induce trust, and how many channels to use. First, one needs to establish the need to use multiple channels, before explaining how to use them.

The problem is that if the Internet is this thing with users and servers attached to its interconnection spigots with nothing but connection pipes between them, and where any attached user or system has protocol-based communication access to all others so mounted, then we must ask what controls the whole thing? And where might we mount a controller for the whole Internet?

The essay explains that the answer depends on how you use the communication concept of trust. You may choose to be at an edge of some local network by joining a mailing list or participating in a message board on some website, or subscribing to some information services, trusting that which has been authorized for you. Or you may choose to be at the center of your own network where you control the nature of all the connections. [p. 19]

Gerck: Internet protocols are much clos-

er to human communication than to Shannon's idealized communication systems. We, speakers of the same natural language, communicate with one another by trading contents, not by exchanging un-interpreted strings of symbols (bits). Each bit of information sent by a human to another must either contribute to the content or be discarded. Content may have different meanings, some of them intertwined, at different layers of our understanding. And, in the same way that content must be conveyed in human-to-human (H2H) communication, we find that content must also be conveyed at different protocol layers in Internet machine-to-machine (IM2M) communication – not just bits. In internetworking, machines are not just trading un-interpreted strings of symbols, or bits. They are trading bits and meaning, machine to machine. They are talking. [p. 20]

The problem is not however in the concept of trust by itself, but in using trust quantifiers that are either artificial or limited in the context of communication systems. A common limitation is to use trust as a synonym for authorization. However, this is valid only in a network, not in an internet. For example, in a network a trusted user is a user authorized by the network management to access some resources. But where is the "Internet management" in the Internet? It does not exist. [p. 21]

The lion does not need to receive any information from the lamb besides that which is communicated in the communication channel itself – the lamb is there. The lamb obviously needs to know whether the lion's assertion "I'm not hungry" is truthful. The truthfulness of this assertion is not information and cannot be transferred using that same channel.

Why not? The truthfulness of the lion's assertion cannot be information because in Information Theory information has nothing to do with knowledge or meaning, and it cannot be transferred using that channel because how could the lamb know that the lion was not lying?

The same situation would apply to two

machines on the Net, or to humans communicating. Information is surprise, and not always nice. Therefore, we see that “trust me” is an empty affirmation; a self-affirmation cannot communicate trust.

In other words, a decision to trust someone, the source of a communication, the name on a certificate, or a record must be based on factors outside the assertion of trustworthiness that the entity makes for itself.

Loosely speaking, we can say that “information is what you do not expect” and “trust is what you know.” In the lamb example, the lamb needs to trust whether the lion is hungry. This could not have been information, because information is what the lamb does not expect – information is surprise. To be sure, the lamb does not want surprises in regard to the lion’s appetite.

Linking both concepts of information and trust, we can say that “trust is qualified reliance on received information.” [p. 22]

Misusing trust is also related to the scenario of a spoofing attack. In spoofing, a user trusts a fraudulent service or information that pretends to be legitimate. Spoofing attacks cannot be prevented by using SSL, digital signatures or encryption. The U.S. National Science Foundation recently sponsored a conference dealing with Internet security issues for voting applications, where spoofing was declared an open and very serious problem without a solution today and long-term research on the subject was recommended.

The problems mentioned above are diverse and touch upon many different aspects of the reliability of Internet communications. We need to solve these problems. A large section of our economy and our lives are already riding on the Internet. [pp. 22 -23]

In control theory, the basic parameter to measure performance is position-error, which translates here to the actual re-

sponse as compared to its expected or estimated (i.e., trusted) response. In open loop-control, one method frequently used to decrease position-error is to introduce periodic checks of any convenient system variable, not necessarily the control variable. This is not a feedback loop because it is done after the actuation. This method is equivalent to the well-known dictum “trust but verify” – implying the need for a pre-defined policy of checks and balances that can periodically adjust the trust estimator as a function of observed behavior.

Thus, trust can also be explicitly defined as “trust is an open-loop control process of an entity’s response on matters of X” or, less precisely but more concisely, as “trust is to rely upon actions at a distance.” [p. 23]

An attempt to bring centralized control to the Internet would also need to deal with that vexing question –what to control and where? In the Internet, no single person can tell which networks are included, because no one is there to tell which networks to include and which not to. Any user can add any number of networks to the system. Centralized control is impossible in an internet made out of open-ended networks of networks.

But if the answer does not lie in centralized control, what is the answer to the problems mentioned above? Could Internet control be decentralized? How would this be effective and not generate even more confusion? How would this provide for planetary reliable interpretation of protocols and their messages by machines? [p. 23]

Indeed, as we observe the real world and seek to model the trust mechanisms that allow business and human communication to function, what do we see? We do not see a hierarchical trust structure controlling business from a single center. We do not see “edge control” either, and where it was tried it resulted in anarchy. What we see are inter-entity (intersubjective, see Glossary) relationships heavily qualified in many ways. We see inter-nets, networks of networks, a manifold

of networks with multiple control boundaries and lacking a common single reference.

Likewise in the Internet, we see a set of edge-edge, center-center (yes, there are many centers), center-edge and edge-center relationships that we can use to induce (communicate) trust. Trust is always formed from relationships between entities because it is induced by, or results from, a communication event. Such relationships may be unilateral (e.g., I do not know you and I do not know that you trust me) but not singular. There can be no information transfer between a sender and itself, since all is known. The net transfer is zero (see sidebar “Trust Points”). The only possibility for an entity to transmit information to itself is to do so to the future, but then if and only if there has been some loss of memory.

Note that so far we are not yet talking about how we use multiple channels to induce trust, as first we need to establish the need to use multiple channels before explaining how to use them.

Second, we need to direct our attention to the last sentence of the sidebar “Trust Points,” quoted from the book “Digital Certificates: Applied Internet Security” by J. Feghhi, J. Feghhi and Peter Williams, where the authors state:

“Models of integrity and trust in a particular space are best left to the communities of interest, governments, and industry forums, which are most familiar with these groups’ specific needs.”

This sentence reminds us that trust is always local to and is earned in communities of interest. This dispels the idea of a cookie-cutter approach to Internet control, since each community of interest (for example, your company) will have different goals, different control objectives. **Increasingly, users want more freedom in controlling their own connections, bandwidth, services and rules of use.** Maybe that’s why NAT/IPv4 is broadly used and IPv6 is not widespread – even though IPv6 is better in many aspects. [p. 24]

Lessig Review, pp. 25-30

[Go to Executive Summary](#)
[Go to Full Interview](#)

Lessig “My central claim throughout [this book] is that there is a benefit to resources held in common and that the Internet is the best evidence of that benefit. As we will see the Internet forms an innovation commons. It forms this commons not just through norms but also through a specific technical architecture. The Net of these norms and this architecture is a space where creativity can flourish. Yet so blind are we to the possible value of a commons that we don’t even notice the commons that the Internet is. And, in turn, this blindness leads us to ignore changes to the norms and architecture of the Net that weaken this commons. There is a tragedy of the commons that we will identify here; it is the tragedy of losing the innovation commons that the Internet is, through the changes that are being rendered on top,” [Lessig, p. 23 and *COOK Report* p. 25 above].

Here is his own description of the task he sets: “Changes threaten the power of those now in power; they will work in turn to protect themselves from the changes. In the balance of this book, I want to detail their work to change the Internet, and the legal culture surrounding it, to better protect themselves. Some of these changes are legal; some are technical; and some use the power of the market’But all are driven by the desire to assure that this revolution doesn’t muck things up-for them. There’s nothing immoral in this desire. This is not a battle between good and evil. Stockholders demand that management maximize its income; we shouldn’t expect management to do anything different. But, even if this is “only business” to them, this does not mean it should be “just business” for us. We need not stand by idly as the Internet is changed. They have their interests; we have ours. And for those who believe that the environment of creativity that the Internet produced was worth something, there is reason to

resist the changes that I will describe.” [Lessig, p. 146].

Lessig says what regulators here should do. “[R]egulators should begin to evaluate changes to the network in terms of the neutrality of end-to-end.” True. But he also seems almost mystified that they have not done this. He states what regulators and other policy makers have just done in Canada. “The Canadians have required open access for broadband providers and a recent Report of the National Broadband Task Force has endorsed a “bill of rights” for broadband users that assures continued consumer choice for any build-out of a network employing government funds. This choice of policy has apparently not harmed Canadian access. According to a recent OECD report, broadband connections are twice as common in Canada per capita as in the United States,” [p. 249]. We conclude that the Canadians have not yet sacrificed all ideas of Canadian national and public interest on the altar of the unfettered free market. Won’t it be interesting if what is derisively referred to as “socialist leanings” in many US political circles creates an environment in Canada that is a fertile seedbed for new technology development?

“A policy from the FCC that does not create a strong opportunity for an alternative to develop is designed to protect existing interests. To not encourage or permit wide-scale experimentation, to not set aside much broader unlicensed spectrum, to protect existing uses against any interference these are policies designed to preserve the old against the new. They are just what we would expect from government regulation of spectrum; they are much less than we should demand after the experience of the Internet. The government’s role should be to induce investment where there is a great deal of social value to be created. This is precisely the opportunity with unlicensed spectrum,”[Lessig; pp. 226-27]. It seems to us that the goals of CaNet*4 also fall into this category. [p. 27]

He implies that if we were educated enough and understood what we were

losing that we would fight against the losses. Therefore, he sets out to educate us.

“What’s at stake here are two models for organizing a communications network, and the choice for us is which model will prevail. On the one hand, there is the model of the perfectly controlled cable provider -owning and controlling the physical, logical, and content layers of its network. On the other hand, there is the model of the Internet-which exerts no control over a physical layer beyond the decision to include equipment or not, and which enables the free exchange of content over a code layer that remains open.”

“As the Internet moves from the telephone wires to cable, who should govern? When you buy a book from Amazon.com, you don’t expect AOL to demand a cut. When you run a search at Yahoo!, you don’t expect your MSN network to slow down anti-Microsoft sites. You don’t expect that because the norm of neutrality on the Internet is so strong. Providers provide access to a network that is neutral. That’s the essence of what the Internet means.” [p. 28]

“As commentator Gordon Cook writes: The Internet revolution has come and gone. It has created a tremendous burst of innovation [-a] burst that now looks to have been mismanaged.... [T]he people who did the least to advance the new technologies seem most likely to control them. We are left not with the edge-controlled intelligence of the [end-to-end] network but with the central authoritarian control of the likes of AOL Time Warner.” [He cites our “What Good is the Stupid Network. . .” article from the September 2001 issue.]

“The irony astounds. We win the political struggle against state control so as to re-entrench control in the name of the market. We fight battles in the name of free speech, only to have those tools turned over to the arsenal of those who would control speech. We defend the ideal of property and then forget its limits, and extend its reach to a space none of our Founders would ever have imagined.”

“We move through this moment of an architecture of innovation to, once again, embrace an architecture of control -without noticing, without resistance, without so much as a question. Those threatened by this technology of freedom have learned how to turn the technology off. The switch is now being thrown. And we do nothing about it,” [Lessig, pp. 267-68 and *COOK Report*, p. 28]

“Our future is this: the free speech clause of the first amendment will be read to entitle those who own the wires to change the logical layer and make it owned as well; the free competition principle of the Sherman Act will be read (by the same circuit we might notice) to entitle the owner of the platform that most affects this logical layer (that one company whose name I have not uttered) to code that platform to discriminate as it wants; and the free culture that we have seen

flourish in this commons built by the Internet will be captured and controlled again by those who control most of the content, and by those who succeed in congress in expanding their control from the imperfect to the perfect.”

“The future of control will get built by an idea; the idea that property is good so more property is better. It will get sanctioned by a culture that has forgotten any distinction, and that is so blinded by what it has forgotten that it does not even notice when the most extraordinary innovation that our culture has seen since Thoreau was a name most Americans could spell is built not on an architecture of perfect freedom; not in a world where every layer is in the commons; but also not on an architecture where control was the rule; not on an architecture where every layer was owned: but instead on an architecture that mixed freedom and control;

that built property within a commons; that got its life from this mix of property and the commons.” [p. 189]

“We allow these changes, they don't just happen. We stand back as they occur, they don't happen in the night. We let them occur because most of us believe they should; control is good, better control is better, these systems of control are ways to make sure the better comes from the good. It is an attitude and blindness and a pathetic resignation that permits this change. So enamored we are with the invisible hand, so convinced we are of the genius of property, so blind we are to what makes innovation possible, that we allow the undoing of the most significant chance for something different that we have ever seen.” [p. 190 - preceding three paragraphs from Duke University Law School Symposium November 9 -11, 2001]

Executive Summary:

Introduction:

While we still have a global Internet and probably always will -- (at least in the sense of something defined as mail and web servers connected via DNS to the ICANN legacy root), it is now clear that technology is taking us far beyond this initial two dimensional Internet. It is giving us the ability to establish our own networks that may or may not be directly attached to the global internet.

We may now network all manner of devices running TCP/IP. As a result of this new ability, we will begin to create new, locally originating networks that grow and shrink on a dynamic ever changing basis in sync with the needs of those behind them. With access to wireless spread spectrum 802.11b devices and before long with access to wavelengths that are becoming more and more affordable, we will increasingly begin to build own physical networks.

However, with other technologies located at the content layer of the network, we can also use TCP/IP to set up and tear

down inter-networked subsets of communications, along with communities of interests and content based on those interests. For some of us the global legacy Internet is now about to become a myriad of tribalized mini-internets. Many people may begin to drift in and out of these communities, using them like the UseNet of old to connect to others with whom they wish to do business. While the legacy Internet will capture the vast majority of available 'eyeballs' for the foreseeable future, we predict that many people will begin to spend portions of their time attached to all of these networks.

Larry Lessig bemoans (in his new book reviewed in this issue) the corporate propelled movement to "enclose" the internet with wall gardens and other restriction of expansively defined intellectual property. While this movement continues to press forward, there will be groups of people unable to do what they want within the realms of the old technology who will begin to experiment with new ways of communication.

CA*net 4, pp. 1-18

Highlights - Full article

The Canadians have a vision that has been lost in the free market purity of the United States' political environment. Consider the theme of the CANARIE's 7th Annual Advanced Networks Workshop being held this week in Toronto.

"Following the recent release of the National Broadband Task Force report, there is increased awareness that a national broadband infrastructure serving all Canadian communities will be critical to Canada's ability to innovate. Information technology infrastructure will be one of the most important vehicles for promoting innovation and improving Canada's productivity, leading to increased wealth and economic growth.

Community broadband networks, provincial networking initiatives and national research backbone networks, are all part of the same continuum of providing a national innovation infrastructure.

In the future, research, education and in-

novation will not be solely a product of universities and research centers. A national innovation infrastructure will allow all Canadians in our schools, communities and businesses, no matter how remote or how distant, to be full participants in developing and using innovative applications and services.

New concepts involving "grids" and "eScience" are coming to assume greater importance in many branches of science. Some of this work could allow students in our schools, and eventually members of the public, to participate in basic research that otherwise they could only read about, thereby engaging them directly in Canada's "innovation culture." See also : http://www.canarie.ca/advnet/workshop_2001/agenda.html

CA*net 4 is to be built in part on the premise that with huge amounts of fiber laid and very large amounts lit but still not fully utilized direct control of actual bandwidth can, for the first time, be placed into the hands of customers and then end users. According to Bill St. Arnaud: "Today networking is like computing was 40 years ago when the market was dominated by large mainframe computers. But in the 1970s the mini-computer came along followed by the PC which fundamentally changed our thinking of how to do computing. Computing became personal. The user was empowered to develop new applications and services that were not possible on a mainframe computer. With CA*net 4 we hope to move networking in the same direction as computing has gone in the last 30 years." Arnaud wants to turn the network itself into a customer owned and controlled asset.

"From day one we will be assigning ownership and control of individual wavelengths or STS channels to the GigaPOPs, universities and perhaps even individual researchers. They will be free to trade and swap amongst themselves and do what ever they want with those wavelengths. From day one we will also encourage these organizations to directly peer with each other and other international research networks using these wavelengths. But, initially the BGP opti-

cal peering will be done manually. Once OBGp is successfully implemented, it will allow these organizations to automatically change the routing of the wavelengths and peering relationships without first contacting CANARIE. So rather than operating a traditional hierarchical IP network as many other research networks do today, CANARIE will only offer an aggregate IP network as an optional service for those organizations that don't need their own wavelengths."

Arnaud and Wade Hong who has done the technical proof of concept development with OBGp over the past year explain in detail how because of the difficulty of making a new protocol within the IETF, they have decide to follow the injunction of SUN's Bill Joy to provision "services by agent exchange rather than by defining new protocols." Bill Joy in his article in *Internet Computing* two years ago was not thinking of OBGp when he wrote: "This change will mark the beginning of the end of the Internet as a world of protocols and the beginning of the era of the agent-based Internet." Nevertheless, what CA*net 4 intends to do could be the first major step in the direction that Joy foresaw.

The ultimate goal is for a researcher to be able to use a web based interface on a workstation to do point and click peering of a lambda between his institution and another institution at a switch located at an exchange point where customers (regional networks, institutions or even end users) own wavelengths and ports. CA*net 4 wants to seed a market place where local networks and research institutions can become customers for wavelengths that many next generation fiber players now want desperately to sell.

As St. Arnaud says: "Rather than having a few hundred IXs for a small number of big ISPs, we will strive toward the point of having thousands of IXs for individual institutions and users. Ultimately, as an end user, rather than depending on your ISP to interconnect you to somebody else, the intent is for you to be able to cross connect to that person just like the two of you were actually interconnecting at an exchange point."

When we asked his thoughts about the ways in which CA*net 4 might help the Internet build a viable business model for life after the current down turn, he replied: "I think the business and architecture model of the future will be of control and management moving increasingly closer to the edge, not only of the in terms of applications, but also in terms of control over the infrastructure. I think that one of the drivers for this will be as a consequence of the issues that of Larry Lessig has raised where content and distribution companies are trying to exert control over the Internet infrastructure to protect their intellectual property interests. Decentralization and minimizing control at the center will help thwart these challenges. We are already working on concepts with our industry and research partners to extend this concept of customer control of wavelengths all the way to the individual home. . . .

In the future we see a physical network infrastructure that closely parallels Morpheus and other peer to peer networking paradigms. The end user will have a choice of whether they wish to subscribe any number of "walled garden" service providers. Or they may chose to physically connect to community networks to share files and data with high speed Gigabit wavelengths bypassing all traditional hierarchical service providers.

A lot of these concepts are still very speculative and unproven at this point in time. But what it does point out - is the critical role that research and education networks still play in the ongoing development of the Internet. . . . Most importantly research and education networks will continue to play an absolutely critical role in exploring new concepts in networking which initially may appear very radical to the traditional telecom world.

**Trust as Qualified . . . ,
pp. 19 -24**
[Highlights - Full article](#)

The trust definition and issues raised in Ed Gerck's essay have been discussed online since 1997 in several technical

groups including the MCG and IETF's PKIX lists, and have also been presented in books and essays. Detailed and in-depth online discussions among the experts, as well as practical applications, have helped Gerck evolve and test these concepts over time. What has been missing is a summary and clarification of the arguments. This is what we present in the article on pages 19 - 24 of this issue.

Trust is a word that is commonly applied to many situations and consequently has many shades of meaning. This essay by Ed Gerck focuses on one precise set of coherent meanings: the concept of trust in the context of communication. More specifically, in the context of the engineering problem of Internet communications. Gerck defines trust as "that which is essential to a communication channel but cannot be transferred from a source to a destination using that channel." Thus, trust is considered something essentially communicable, but with specific rules for its communication. For example, self-assertions cannot induce trust. Client-server communication is not enough to induce trust. Gerck demonstrates why trust is needed and shows the interplay between trust and power. His exposition also discusses the induction (communication) of trust in heterogeneous environments, from human to machine, machine to machine, and machine to human.

Note: Meaning must also be communicated in such heterogeneous environments. In fact, introducing meaning into information theory so that the communication of meaning can be described, has been an open problem since 1948. Gerck asserts that the way to communicate meaning is to first communicate trust and bits, and then use them to define the meaning. In other words, rather than introduce meaning into information theory, we introduce trust. Meaning will be introduced and conveyed implicitly.

The problems of trust are diverse and touch upon many different aspects of the reliability of Internet communications. We need to solve these problems. A common question is whether these problems could not be solved by more con-

trol. In the discussion of trust versus control, Gerck notes that it is instructive to view trust as an open-loop control process, i.e., a control process that does not rely on a closed feedback loop in order to achieve its purpose. This approach to trust in communication systems versus close surveillance has a number of advantages, including simpler systems, easier design and interfacing, less cost, and the lack of a single handle of control which is, of course, a single point of failure.

While general issues of control on the Internet have been talked about for many years, the specific issue of bringing centralized control to the Internet would need to deal with that vexing question-- what to control and where? Gerck reminds us that in answering this question one must first dispel the notion of "center" and "edges" existing in the Internet. The Internet is not a network; it is a network of networks. Talking about "edges of the Net" is like talking about the "last Web page of the Net"-- where is it? This line of reasoning also shows that the argument for "trust at the edges of the Net" is a fallacy. An edge that is able to route IP packets becomes a center.

Furthermore, Gerck states that we need to remember that trust is always local to and is earned in communities of interest. This speaks against the idea of a cookie-cutter approach to Internet control, since each community of interest will have different goals, different control objectives. In summary, the answer needed to solve the fundamental problem of Internet communications is trust. Not trust as blind faith, compliance, belief, or ignorance, but trust as qualified reliance on information through open-loop control. Trust is that which provides meaning to information. Trust is something that can flow in our existing TCP/IP, dial-up and other networks.

Lessig: Future of Ideas, pp. 25 -30
[Highlights - Full article](#)

Lessig's book explains why and how corporate America has brought the seemingly uncontrollable Internet under its con-

trol. It was too successful and too ripe a target for its own good. Turned commercial and subject to tremendous hype as the engine of and means for commerce for the 21st century, the Internet quickly became a strategic corporate target for the telecom, computing and media and content sectors. It was seen as the next great source of wealth and power. As such it had to be exploited and controlled.

Lessig "My central claim throughout [this book] is that there is a benefit to resources held in common and that the Internet is the best evidence of that benefit. As we will see the Internet forms an innovation commons. It forms this commons not just through norms but also through a specific technical architecture. The Net of these norms and this architecture is a space where creativity can flourish. Yet so blind are we to the possible value of a commons that we don't even notice the commons that the Internet is. And, in turn, this blindness leads us to ignore changes to the norms and architecture of the Net that weaken this commons. There is a tragedy of the commons that we will identify here; it is the tragedy of losing the innovation commons that the Internet is, through the changes that are being rendered on top," [Lessig, p. 23]

Here is his own description of the task he sets: "Changes threaten the power of those now in power; they will work in turn to protect themselves from the changes. In the balance of this book, I want to detail their work to change the Internet, and the legal culture surrounding it, to better protect themselves. Some of these changes are legal; some are technical; and some use the power of the market' But all are driven by the desire to assure that this revolution doesn't muck things up-for them. There's nothing immoral in this desire. This is not a battle between good and evil. Stockholders demand that management maximize its income; we shouldn't expect management to do anything different. But, even if this is "only business" to them, this does not mean it should be "just business" for us. [p. 146.]

Lessig summarizes with scholarly detail the war waged by the interests of corporate control on behalf of corporate profits against the physical, logical and content layers of the Internet. His conclusions are steeped in pessimism and anger. In our opinion he expressed them even more forcefully in a paper at a November 9-11 2001, Duke University Law School symposium. "If communism vs. capitalism was the struggle of the 20th century, then control vs. freedom will be the debate of the 21st century. If our question then was how best to control, our question now will become whether to control. What would a free resource give us that controlled resources don't? What is the value in avoiding systems of control?" [p. 178]

"The future of control will get built by an idea; the idea that property is good so more property is better. It will get sanctioned by a culture that has forgotten any

distinction, and that is so blinded by what it has forgotten that it does not even notice when the most extraordinary innovation that our culture has seen since Thoreau was a name most Americans could spell is built not on an architecture of perfect freedom; not in a world where every layer is in the commons; but also not on an architecture where control was the rule; not on an architecture where every layer was owned; but instead on an architecture that mixed freedom and control; that built property within a commons; that got its life from this mix of property and the commons." [p. 189]

"We allow these changes, they don't just happen. We stand back as they occur, they don't happen in the night. We let them occur because most of us believe they should; control is good, better control is better, these systems of control are ways to make sure the better comes from the good. It is an attitude and blindness

and a pathetic resignation that permits this change. So enamored we are with the invisible hand, so convinced we are of the genius of property, so blind we are to what makes innovation possible, that we allow the undoing of the most significant chance for something different that we have ever seen." [p. 190]

ICANN's Broken Promises, pp. 31-32

[Full article](#)

An article from ICB Toll Free news summarizes the latest series of contortions by ICANN's leaders to explain why ICANN's secretive intellectual property cadre cannot afford to keep promises made to Congress by the likes of Esther Dyson and others. The sad and dismal on going story of the self-selected elite that pretends it still has a mission critical to the stability of the internet.

The United Kingdom Has its Own Local Loop Problem

On November 21, 2001 BT's national IP backbone melted down and **Sean Donelan** wrote to the NANOG list: "The joy of single provider service."

From Denmark **Sean Doran** responded: "No, you mean: the joy of a rapacious monopoly with a spineless regulator."

"BT has successfully stalled opening itself up to competition as required by EU and UK law, by sitting on the local loop. Meanwhile, they themselves are deliberately slow on rolling out modern broadband services (e.g., xDSL) compared to the EU average."

In short, the UK is *the* backwater of Europe when it comes to high-speed Internet connectivity -- it is rare to find at all, and when you find it, it's not cheap. That BT's outage affected so many people and organizations is a horrific indictment of the disastrous policy line the regulator has taken, and the incompetence of the government when it

comes to promoting its stated aim of increased connectivity for UK residents.

The official line is at:

http://www.oftel.gov.uk/publications/local_loop/llufacts/llufacts0501.htm

Even the OFFICIAL statistics are depressing:

Number of fixed lines in the UK: 35 million
Number of BT ADSL connections: 50 thousand
There are FOUR sites where local loop unbundling can happen, and they are not remotely in the biggest population centres.

By comparison:

<http://news.zdnet.co.uk/story/0,,t269-s2092106,00.html>

BT admits that it has only handed over 163 residential lines to other operators

By comparison, during last night's election coverage in Denmark, there were 29000 DSL-speed connections to streaming servers operated by one of the main TV operators. Nearly everyone I know in the IT field has DSL at home. Operators are offering low-price SDSL.

Unbundling is working.

Not only is Denmark a much smaller economy than the UK, it also has some very different geographical challenges: it is made up of 400 islands. Yet, even in areas of relatively low population concentration (esp compared to, for example, London), it is possible to get DSL connectivity provided from at least one operator, and often several (4+).

Most EU countries can claim the same success as DK.

Shame on the UK's government, which is directly responsible for the regulatory environment which deprives a huge percentage of the population from getting better, cheaper service than they get from their expensive, unreliable monopoly called British Telecom!

Donelan: [BT lost] of its national IP backbone on Tuesday affecting DSL and Dialup service across multiple ISPs in the UK.

Doran: These ISPs had NO CHOICE with regard to multihoming. They can thank their government.

Donelan: According to news reports,

this affected almost all DSL service in the UK.

Doran: Right, because of the 50000 or so DSLcustomers in the country, a maximum of *163* are not supplied by BT.

Donelan: What happened to BT? Is this a unique "feature" of the UK marketplace or can the same thing happen in the USA?

Doran: You have evil incumbents in the USA too. Beware.

Donelan on November 23: BT is telling ISPs the reason for the multi-hour outage was a software bug in the interface cards used in BT's core network. BT installed a new version of the software. When that didn't fix the problem, they fell back to a previous version of the software.

BT didn't identify the vendor, but BT is identified as a "Cisco Powered Network(tm)." Non-BT folks believe the problem was with GSR interface cards. I can't independently confirm it.

Editor: when we asked a subscriber in England to comment we received the following reply.

I would agree with the comments about xDSL - no matter who you buy it from it is BT Excite who install it. I don't live that far from the nearest enabled exchange (about 1.5 miles), but 3 attempts including the extended reach service and numerous visits by engineers and they could not get anything to work properly.

There is an alternative - cable modem. NTL passes I believe ~6M homes (have a look at www.ntl.co.uk) and offer both a residential and business cable modem service. This what I hope to get installed - though NTL are reknowned for their poor customer service!?!

I am desparate for BB, but its is proving very difficult to get.

New Features: To find out how to use new features, click [here](#) with Acrobat reader hand tool while connected to the internet Or go to <http://www.cookreport.com/features.shtml>

Betting Your Company's Future on Your Favorite Business Model: "Empowering the Customer" or "Empowering The Telco?" to be published early January 2002 at \$385

Subscription Rates

Choice of either ascii or Adobe Acrobat (PDF) format 1. Individual; College or University Department; or Library; or Small Corporation - \$295 2. Corporate - (revenues \$10 to 200 million a year) - \$375 3. Large Corporate- Revenues of \$200 million to \$2 billion per year - \$475 4. Very Large Corporate- Revenues of more than \$2 billion per year - \$575 Site License: The right to distribute ascii and PDF via email to all employees of corporation. 5. Small corporate: \$450 6. Corporate: \$700 7. Large Corporate: \$950 8. Very Large Corporate: \$1200 . Site License Distribution via intranet web site \$400 a year additional. See www.cookreport.com for more detail

Gordon Cook, President
COOK Network Consultants
431 Greenway Ave
Ewing, NJ 08618, USA
Telephone & fax (609) 882-2572

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