



Gig and Ten Gig Ethernet Bypassing Carriers and Revolutionizing Transport at Internet's Edges

Gigabit Ethernet Combined with Metropolitan Area Fiber Enables New Data Network Business Model in Montreal

Editor's Introduction: Gigabit and now ten gigabit Ethernet is bringing revolutionary change to provisioning of IP bandwidth at the edges of the Internet. The result enables the complete elimination of SONET from data networks. It also enables individual companies to eliminate carrier managed circuits by making it economically attractive to use Ethernet across both wide and metropolitan networks of leased dark fiber without relying on carrier services. In two interviews in this issue we track the effects. In the first Francois Menard describes how gigabit Ethernet is changing the business model of providing data network services in Montreal. In the second Cisco's Howard Frazier describes ten gigabit standards and explains the economic underpinning behind gig and ten gigabit corporate private data networks.

Menard has appeared in the pages of the COOK Report on several previous occasions. He is currently Project Lead, Internet Division of Cable VDN Inc. in Montreal, Canada.

COOK Report: Is it your position that there are two possible approaches that can be taken in building data networks at this point in time? One of them is the VDN and Cogent Communications approach of throwing bandwidth at the user - not that they're the only two companies in the world doing it - while the other would involve a lot of MPLS and circuit-switching types of things at Layer 2. If one is trying to paint the big picture, is this a reasonably accurate summation of the big divisions?

Menard: Yes, but to complete the statement, you should say that the two different views aren't just bandwidth and packet-switching versus circuit-switching or MPLS or label-switching, for that matter. They are two different world-views, one where people build

networks to run certain applications and the other where people build applications to run on certain kinds of networks. If you're in the camp that builds applications for networks, then you end up building applications that always maximize the use of the network, rather than a network that maximizes the use of a certain set of applications.

The applications that maximize the use of a network are traditional IP, and adaptive applications, which rely on the Transmission Control Protocol to put out more features when bandwidth is available and which peacefully back off when there is congestion.

COOK Report: Let's go down that pathway then in some detail, if we can. The basic insights of Bellheads versus Netheads <<http://www.tmdenton.com/netheads.htm>> have been here for a while and haven't much changed in the last year or so. What's been changing very rapidly are the productivity increases in fiber optics and lasers and optical networks and the increases in the bandwidth that you can get out of a basic piece of optical infrastructure, especially now with gigabit Ethernet available in the local loop. From that point of view, how is this changing the way in which you will find it most cost effective to build IP networks? Please illustrate the changes with some examples from your own company.

Menard: Basically, it doesn't cost us very much more to deliver to a customer ten megabits, one hundred megabits or even gigabit, because we essentially dedicate a pair of optical fibers point-to-point to every location whom we serve.

COOK Report: And from that point of view, does it really become feasible to throw bandwidth at a problem for the first time?

Volume IX, No. 7, October 2000
ISSN 1071 - 6327

Menard: Yes, because the transceivers are cheap enough. Provided that you are in a situation like we are in Montreal where the city utility owns the conduits in such quantity that it's not a problem to make a point-to-point connection between two big, shiny buildings in a downtown environment or between buildings even further out in the neighborhoods

Because of the ease of access to fiber the wireless companies entering the market do not have a large advantage over the players that are bringing wire line networks into Montreal. The time needed to deploy a wire line infrastructure, which would currently be a fiber optic infrastructure, is quite short, providing you have the money and the license to deploy. I think almost at some point the bottleneck becomes the availability of the optical fiber itself, because there's so much of a big backlog on the availability of fiber today. Monopoly incumbents place a tremendous stress on availability of fiber for new entrants like us.

COOK Report: Do you mean just getting the cables from Owens-Corning or whom-ever?

On the Inside:

Gig E & Net Design	pp 1 - 9
Economics of 10 Gig E	pp. 10 - 15
Ericsson on Wireless	pp. 16 - 21
IPv6 Allocation & cost	pp.15, 21-22
Executive Summary	pp. 23, 24

Menard: Yes. From Siecor, Perelli, Lucent, Hitachi, all of them.

COOK Report: So with all the conduits and the fiber there, can you give me some idea what your company does own by way of dark fiber?

Menard: We have about 200 route kilometers of fiber.

Economics of Point to Point Fiber in a Metropolitan Area Network (MAN)

COOK Report: What about the pricing issues? How do you decide on how this is priced?

Menard: The pricing policy of VDN is a dollar per meter per month for a five-year commitment with discounts available for longer terms and high volumes..

COOK Report: And that works out to how many dollars per kilometer?

Menard: It is a thousand dollars a kilometer a month for a pair of fibers. But the point is that fiber itself is reasonably cheap. There is every reason to expect that by being able to throw some small, Wavelength Division Multiplexing (WDM) equipment at each end, it would be very logical to price gigabit Ethernet service over fiber at no more than a dollar a meter a month. If you do this, you end up with something that has a lot more capacity than the traditional SONET-based T3 kind of service at a fraction of the price.

Now rather than selling dark fiber for a thousand dollars a kilometer per month, my inclination becomes to sell gigabit Ethernet for a thousand dollars per kilometer per month or less. For if I were to sell dark fiber or lambdas to a competitor at a thousand dollars a kilometer, that competitor could challenge on a single of my strands a good portion of all of the additional capacity available on the other strands.

COOK Report: How does that work? We know that there's a lot of pretty powerful, expensive equipment that can put a lot of different wavelengths on a strand. What about some low-end equipment? Has that become available recently?

Menard: Yep, it's definitely becoming available. Soon there will be offerings in the WDM metro space from Nortel (through its acquisition of Cambrian Systems in October 1998), Cisco (through its recent acquisition of Qeyton Systems in May 2000) and Lucent (through its acquisition of Chromatis Networks in June 2000). Nbase* Xyplex, is

a subsidiary of MRV Communications, www.mrv.com. Nbase which makes the Gigabit Ethernet (GigE) transceivers used in a lot of traditional boxes, like Cisco switches, also makes a Coarse WDM (CWDM) quad GigE multiplexing equipment for a less than \$20,000. That is 4 times GigE on a pair of fibers. Extreme Networks also has a module called WDMi for its GigE switches, which multiplexes 8 GigE channels on a pair of fibers. However, with the scheduled availability of 10 megabits per second transceivers before the end of this year and the projected availability of 100 megabit per second ones, great challenges will face metro WDM equipment at sub 10 GigE speeds.

Furthermore, with all three big players (Cisco, Lucent & Nortel) offering proprietary metro WDM equipment, one can expect standards and interoperability issues to delay the adoption of such equipment. With convergence on a single framing protocol (Ethernet) and the availability of faster transceivers (10GigE), complex and expensive multiprotocol metro WDM equipment fluent in GigE, ESCON, Fiber Channel, and Sonet protocols, faces multiple challenges. The metro WDM equipment needs to decrease rapidly in price, to integrate 10 Gig E instantly and to begin to mux at greater speeds. It needs to do all of this just to continue to offer a compelling value proposition compared to 10 GigE. Although one can only speculate on what will happen with 100 GigE.

COOK Report: Suppose you had a deep bank account and you wanted to go for a high-end piece of equipment that had 32 colors or 40 colors. I have no idea what they cost, do you?

Menard: It's in the \$50K range per set of four colors.

COOK Report: You mean each increment of four colors is an increment of \$50,000?

Menard: Yes.

COOK Report: So if you go to 16, that's \$200,000. And 32 is \$400,000.

Menard: And this is per pair of fibers. Imagine what space will be taken by equipment that will drive a 144 strand cable. It's very expensive.

COOK Report: So the point is, if a smallish company wants to come into the market and buy a fiber pair, such a company can put some significant bandwidth on that fiber pair for a fairly low entry price. Is that what VDN is doing?

Menard: Yes, VDN has both 144 and 288 strand cables that extend to over 200 route kilometers across the city of Montreal. To-

tal strands owned come to over 10,000 kilometers.

COOK Report: That's a lot of fiber.

Menard: The point is, if you're going to string fiber, you might as well string fiber with a lot of capacity. Fiber continues to be installed this summer under the city with cable having well in excess of 200 strands. There's more fiber underneath Montreal today than anybody will ever need in his or her lifetime, that's absolutely crystal clear.

Editor's Note: We did some further research into Montreal's underground utility infrastructure and found out that it dates to the beginning of the 20th century where the City of Montreal made a deal with Hydro Quebec, the province's nationalized electrical utility. Many years later with the creation of the subway and the staging of the 1976 Olympics, the city said that it would create and then manage the infrastructure for bringing communications cables into the basements of buildings. After all it had done basically the same thing for electricity.

Compared to some other cities, Montreal is a very underground city. You can walk for kilometers under Montreal and do a lot of underground shopping. Most of this underground infrastructure comes from 1976 Olympics and the construction of the subway network. The city of Montreal decided that it wanted to prevent further digging up of its streets. Consequently, it put in a lot of conduits. The entity within the City that accomplished this is called the CSEVM or Commission des Services Electriques de la Ville de Montreal. In English; Electrical Service Commission of the City of Montreal. The CSEVM conduits run not only through the streets, but also into each and every building. Looking at other Canadian cities only Toronto has some of the installed underground fiber capability that is to be found in Montreal.

COOK Report: To go back then to the issue of VDN's 144 and 288 strand fiber cables, it actually put those cables underneath Montreal in, what, the last year or two?

Menard: Yes, VDN was the first company in Canada to get a cable TV distribution license in the Montreal market, which was already being served by an incumbent cable television carrier. Incidentally, the timing was very good, just before the CLEC explosion in Montreal. Now we have seven CLECs in Montreal with Central Offices equipped with real Class 5 switches prepared to sell you POTS services: Videotron Telecom, ATT Canada/Metronet, Axxent/Optel, Telus/Quebecel, Groupe Telecom, Norigen and soon C1 Communications. None of them in the residential sector, of course. One can expect major consolidation

across these competitors in the nearby future.

COOK Report: So VDN may be thought of as kind of the first Canadian cable TV CLEC?

Menard: In Canada, the CRTC calls us a BDU, a Broadcasting Distribution Undertaking.

COOK Report: Are there any figures on what it cost them to put down that 200 kilometers of cable infrastructure? Something like \$10 or \$20 million?

Menard: Sorry I can't comment on that.

COOK Report: The reason you're not commenting is that the figures are proprietary or you just don't happen to know them at the moment?

Menard: If I would know them, they would be proprietary, but I really do not know them. I don't think anybody would reveal the installation cost. Nonetheless, figuring out the cost structure of stringing fiber in Montreal is as easy as finding the prices of Siecior and the rates for the rental of conduits. It's pretty much the price of the fiber itself, splicing it and acquiring the rights of way.

COOK Report: And a certain amount of money to string it through the conduits. Presumably if you're in the industry that builds fiber, you deconstruct the various stages of putting fiber networks down into various component parts. As you just said, if you are digging up X number of kilometers of streets, that's really super expensive and there's an installation cost that adds up, right?

Menard: What you say is true. In Montreal due to the nature of the city's conduit infrastructure, this installation cost you mention is not as large as in other cities

COOK Report: Consequently it's probably more cost effective if you have a lot of fiber in the ground to take the \$15,000 optronic boxes and put, what, four gigs on each strand?

Menard: Well, the logic goes this way. When you can face competition from somebody who rents fiber from you and can throw expensive WDM equipment at the problem of increasing capacity, your most effective cost structure comes from your ability to not have to pay for an expensive WDM infrastructure up front. Having lots of cheap fiber, you grab a new fiber pair and just buy new inexpensive optronics to put on the new pair every time that you run out of capacity.

Your design criteria are to build GigE networks on fiber - networks that are as simple as possible to build.

Economics of Gigabit Ethernet or the Business Model of Infinite Bandwidth

COOK Report: And as inexpensive.

Menard: And as inexpensive as possible. What you end up selling is not lambdas or the fiber itself. You end up selling optical data services. They just happen to be provided on dedicated fiber, but at some point in time, you'll be taking control back of that fiber and be putting that fiber on something like a Nortel Optera Metro box.

So the idea is that you don't give out too many 20-year Irrevocable Rights of Use (IRUs). You tell customers that you are selling them GigE and that you reserve yourself the right, at some point in time, to take back control of the strand and install WDM equipment in order to get more out of the fibers already in use. At this time, the price of WDM equipment will hopefully have fallen significantly. So your competitive advantage comes from being able to defer the cost of the WDM equipment as far as possible into the future.

In Montreal, while we are planning to commercialize our network as optical data services rather than as IRU's on lambdas or on fiber strands, we are seeing that it may not be enough to satisfy certain customers. In the last two months, several primary and secondary school commissions in the province of Quebec have commissioned, through open tenders, dark fiber networks to link all of their establishments. It is to be expected that two years from now, in Montreal, schools without access to their own dark fibers will be the exception rather than the norm.

You can view the list of establishments at the following URLs: WWW.CSDM.QC.CA (>200 schools), WWW.CSDGS.QC.CA (>50 schools), WWW.CSMB.QC.CA (>80 schools). We can only view the decisions from the provincial government to permit that the tax dollars of the population be allocated to the building out of private optical network infrastructures as a preemptive strike against the commercialization of optical fiber networks as optical data services rather than as IRUs on lambdas or on fiber strands. The ability of carriers like VDN to sell GigE optical data services instead of dark-fiber or lambdas is even a bit questionable. The challenge of established infrastructure suppliers or new entrants such as VDN is no different than before. It amounts to maximizing the revenues that can be extracted from the infrastructure by not selling it away piece by piece, strand by strand or lambda by lambda. Only by supplying

more bandwidth than customers will ever need will the priority shift back from bit transport to applications.

COOK Report: You just raised a lot of interesting issues. Let's go more slowly through how this all adds up. If you put only a gigabit per second on a fiber pair within the city, are you even using one of the \$15,000 boxes that you mentioned a moment ago?

Menard: Not even, we're only using standard Gigabit Ethernet switches. Very ordinary equipment to say the least. The idea is to use simple equipment to provide simple data services at affordable prices.

COOK Report: And what do those guys cost?

Menard: A 6509 from Cisco Systems, with 10 kilometer lasers costs less than \$2000 per GigE port. Companies like World Wide Packets are expected to release Ethernet switches for the home with one GigE input, 8 fast Ethernet outputs and 4 telephone jacks for less than one thousand dollars, and this before the end of this year. This kind of generic equipment is expected to become generally available from several equipment manufacturers, including Asian manufacturers such as D-LINK Networks, which has a fairly extensive GigE switch product line as well.

COOK Report: Therefore one of those boxes from point A, say a major telecommunications hub in Montreal, could be a hub or head end for 16 different one gigabit incoming fibers from other parts of the city?

Menard: Yes. The current VDN architecture revolves around Head-ends equipped with several GigE switches connected to several other buildings called hubs. These hubs are connected to a last wave of buildings where the signal is broken down from GigE to Fast Ethernet over copper to the customers. Ethernet channel bonding (i.e. Cisco Etherchannel) permits the network operator to install additional bandwidth very quickly where aggregated capacity in the uplink causes too much collisions on the downlinks.

COOK Report: To do the cable TV end of your business, you're putting out hubs around the city anyway?

Menard: Yes, correct. So we essentially end up putting a GigE switch into each of these hubs.

COOK Report: That same 16-port head end one?

Menard: In this case, it's a lower-density

GigE-in, GigE-out switch. So it's more like a 3508XL from Cisco. Most Ethernet switch manufacturers have this type of products today.

COOK Report: This smaller switch costs about what?

Menard: Fully populated with long-reach single mode fiber laser interfaces, it usually costs less than \$2500 per port because of the lower-density. However, sometimes it is worth paying a little bit more to use external media converters which are now becoming available in rack-mount configuration from companies like Lancast or Nbase. While this technology is frequently used in Fast-Ethernet configurations, it is slow to penetrate the GigE market due to the delay in the availability of standardized copper GigE interfaces on PCI network cards and on stand-alone Ethernet switches. Soon however, when copper GigE interfaces become readily available, the use of external media converters at GigE speeds will create another price drop in the cost per port of GigE.

Finally, you may create a point of demarcation between the optical portion of your network and the electronics portion of your network, which is the portion of the network that contains all the Ethernet switching gear. Such a boundary is very appealing from a network operational point of view. Instead of putting lasers into the Ethernet switches and hooking each of these switches together across the core of the network, you build the core of the network using boxes with laser output on one side, while, on the other side uses multi-mode fiber light-emitting diodes or copper Ethernet interfaces.

While the cost of multi-mode fiber interfaces are much cheaper than the single-mode fiber interfaces today, companies like World Wide Packets are probably betting on significantly reducing the cost of single mode fiber transceivers to prices affordable for the mass market.

COOK Report: Is this a trick that you figured out?

Menard: It's a trick that Nbase pushes. What they claim you can do with this "trick" is to lower the cost of the whole total network by about 20%. This also gives you an easier and cheaper transition path towards a WDM environment, as single mode laser interfaces are not cost effective when required simply as an interface to nearby WDM equipment

COOK Report: Do you get your redundancy in the routers rather than from any SONET types of things?

Menard: Absolutely. While WDM equipment forgo the need for SONET technology by providing redundancy in the form of

path diversity at layer two, routers can still quickly recover from unavailable routes by re-computing routing tables dynamically. This typically takes a little bit longer, but since TCP/IP applications are much more tolerant to lost packets, the savings brought forward by managing redundancy at layer 3 are usually worth the additional complexity. Ethernet channel bonding over physically diverse paths also allows redundancy at layer two in a GigE environment while not resorting to the use of SONET or recomputing routes at Layer 3.

COOK Report: And you want to stay way away from SONET because of its cost.

Menard: Right. Since we are building the network from the ground up beginning the summer of year 2000, the use of SONET equipment was not even considered. While am familiar with the technology, it was an implicit decision to delay the provisioning of SONET services along with the one of multiprotocol WDM metro equipment. By essentially focusing all of our attention to the provisioning of high-quality GigE services over optical fiber, VDN did not have to consider hiring SONET-fluent engineers, which translates to additional savings that can be passed on to its customers.

The use of SONET simply to provide path diversity was deemed unnecessary since this function is either built into metro WDM equipment, or it can be engineered at layer 3 by dynamically recomputing routing tables. It may also be made available at layer 2 in an Ethernet environment by using Ethernet channel bonding over physically diverse paths. For example, the Nortel Optera Metro box is built this way, with one path going east and one path going west and if one of the two fails, all traffic is transferred on the other path.

COOK Report: So it's like two lanes of a highway and you can get back to the box by going either direction. You eventually go around in something vaguely like a circle?

Menard: Yes. But path diversity using metro WDM equipment is still more complex to plan for properly than redundancy at the layer 3 using conventional TCP/IP routing protocols or by using Ethernet channel bonding at layer 2.

COOK Report: That helps. So because VDN's primary purpose is to deliver cable TV, you're laying down, in effect, what can also serve as a really nice data network over the entire city. Consequently you've got many, many different points that you can take a customer which just wants IP data and get his packets from one point on the map to another point.

Menard: Yes, absolutely.

COOK Report: And tell me a little bit more about how that would work. I mean, the IP data customer just piggybacks on the rest of the infrastructure, right?

Menard: Yes. Another important point about the VDN architecture is that 6 pairs of optical fiber are delivered to each and every building that VDN serves. The first pair is dedicated to cable television, the second pair to data and the last pair is reserved for future use. This is where the name VDN takes its significance as it means Video Data Networks, which recognizes the present need for a dedicated network infrastructure for the delivery of video programming.

COOK Report: And what are your thoughts about pricing links like that?

Menard: The idea is to price metropolitan optical data services more attractively than the rental of lambda's or fiber strands, and this means less than a dollar per month per kilometer for a five-year commitment. While the final pricing scheme has not been set yet, our network will enable VDN customers to receive an Ethernet VLAN type of service between two points in the city at very attractive prices, but also to have access to other physical layer protocols.

COOK Report: And when you say VLAN, you mean virtual private local area network?

Menard: Yes. When you talk about a VLAN, usually you also mean something that's multi-protocol and not just for TCP/IP. And, again, what people sometimes tend to forget is that when you're providing optical data services, there are many other optical protocols out there for which people are prepared to pay money, not just Ethernet, but such as FDDI even fiber channel or ESCON. With these optical protocols, a computer in the eastern part of the city can talk to a hard drive physically located in the western part of the city, enabling physical redundancy and fault tolerance over great distances.

COOK Report: And some of the other possibilities?

Menard: SONET emulation on WDM for people that have zillions of dollars invested in their equipment that don't expect to have GigE interfaces or other interfaces. And you've got more traditional type of fiber optic interfaces, like FDDI (the Fiber Distributed Data Interface), which have been used for a little while to do high speed data transfer between computers.

COOK Report: Are all of these really non-IP mainframe types of things?

Menard: Non-Ethernet/IP applications

would be a good way to put it. If you have network-attached storage systems with Ethernet and TCP/IP interfaces, there is no real reason to resist full migration onto Ethernet and TCP/IP for all applications, including clustering and storage applications.

COOK Report: We are also talking application service providers?

Menard: Yes, banks, and ASP's will want to have physical diversity in a city and have hard drives in one part of the city and servers in another part of the city, and vice versa for the other redundant sites.

COOK Report: Then assume that I'm the Bank of Montreal or some large trading company in Montreal and I want to hook up maybe three different data centers with three other things, like an exchange commodities trading floor. Or assume that I am constructing a hypothetical, large, corporate virtual private network that would run on your fiber in the city. If I'm the guy who's going to start buying this, what would I do? I would give you the points I wanted to reach and I'd tell you whether one gigabit between each point would be enough or whether I need multiples to start?

Menard: Yep.

COOK Report: And so I would start up my applications and we would size things. But what happens if I find that I have hard disk farm in one part of the city and the commodities trading in another part of the city, and, all of a sudden, the volume of commodities trading goes way up? All of a sudden it looks to me like I'm going to hit my data ceiling on your network?

Menard: Well, with GigE, it would be very probable that you would not hit that data ceiling. Your question is a good introduction to the issue of over-provisioning as a service provider in order to sell bandwidth on demand. I personally consider over-provisioning on behalf of your customer, a very poor technical solution to network engineering. If you "over provision," you are making the assumption that your customer is unable to operate his network with adaptive applications.

COOK Report: And if it looks now like I'm going to saturate what I bought next week, what do I do?

Menard: The typical answer from a service provider will be that it has thought about that upfront and has accordingly over-provisioned its network in anticipation of your future needs. But my point of view is that for bandwidth to be provisioned on demand, your network needs to be over-provisioned

and that if your network is over-provisioned, why not make available all of this additional capacity to the customers at all times? Over provisioning mechanisms or Committed Access Rates are nothing but on-demand billing frameworks. And as I have said before, the Nethead idea is to build very inexpensive networks with little to no complexity related to billing and provisioning.

COOK Report: And if I'm a potential customer and we're talking about hypotheticals, what would you tell me I should be thinking of?

The Two Schools of Network Design

Menard: Again, it becomes the two schools of network design. Do you build networks for adaptive applications or you build networks for non-adaptive applications? And most people tend to build networks for non-adaptive applications. Consequently they identify a particular set of applications and optimize the network for these applications.

COOK Report: And this hypothetical user would assume that with the hard disk backup of the commodities trading operation that he probably isn't going to need more than one gigabyte of bandwidth at any one point in time over the next month. What happens when he looks at his traffic figures and he says, Oh, boy, it looks like I'm going to hit ceiling in a couple of weeks?

Menard: Well, I'm assuming that by providing GigE service to the customer, I am providing him so much more bandwidth than he needs that he will not need another GigE link for quite some time. Instead of providing bandwidth in increments of megabits per second, the bandwidth is provided in increments of gigabits per second at a small premium instead of a thousand times more expensive. For as long as this hypothetical customer does not forecast a thousand fold increase in bandwidth needs every month, my assumption holds. But in the case where the customer is indeed forecasting demand in thousand-fold increases monthly, we'd have spare strands that are available out there.

COOK Report: Your strands to whose equipment?

Menard: To the customer's equipment. Again, what I am trying to portray is the strategy that, as a carrier, you can afford to over-provision your network in anticipation of your customer's needs is a poor one.

COOK Report: So what you're saying is you're going to have equipment on strands from his point A to his point B that's sitting there and is probably unused and probably unrented at any particular moment in time?

Menard: Not even equipment, just pure raw fiber capacity that's lying there in the ground being unused.

COOK Report: Oh, well, you've made it very plain you have oodles of that.

Menard: That's the point. Even as a small player, VDN can afford to have a lot of spare capacity. Every time one unit of spare capacity is allocated to one customer, for each of these units, the customer gets at least 100 times more bandwidth than he really can foresee using, and this at prices that are just a touch higher than what he would pay for data network services at that correspond to his exact needs at any point in time. As a small player, I find it simpler to exploit as much as possible the 10x bandwidth increase slope available with Ethernet over optical fiber to put together compelling telecommunications service offerings that leverage Internet adaptive applications to create additional value and thus avoid resorting to sophisticated billing mechanisms based on poor technical concepts such as "committed access rates".

COOK Report: Okay, what can you afford to do?

Menard: So what I can afford to do is have lots of capacity in the ground unused and be ready to go splice that fiber on demand when the customer wants it.

COOK Report: Splice it into pre-existing gigabit Ethernet switches that you have?

Menard: Yes. In this case, into the customer's equipment. The idea of building networks for adaptive applications is that you totally, totally, totally over-provision the customer's network with bandwidth that he does not foresee using and this at prices that are not much higher than building networks that correspond to the exact needs in bandwidth of the customer.

COOK Report: And the customer? Does he think he's being charged excessively?

Menard: Well, the customer understands that the Internet model creates pricing schemes, which allow him to have the flexibility of running adaptive applications and not being charged for sleeping bandwidth.

COOK Report: In other words, you own the fiber and connected to that fiber you are spreading gigabit Ethernet switches all over Montreal. Therefore you can offer a customer a gigabit per second local loop for the price that Bell Canada might want to charge for ten megabits of SONET-protected, managed carrier connectivity? So, if you have an infrastructure that can provide gigabit

connectivity when most folk are thinking they need only ten megabits, you are certainly offering them an over-provisioned network.

Menard: Right. So the notion of bandwidth on demand is service over-provisioning on the basis that you're going to sell network services at speeds which are too closely optimized for what you believe to be the customer's maximum demand for bandwidth.

So let's take an example. I'm a bank. I could tell you that today, my bandwidth requirements are such that I'm pretty sure that 99.8% of the time I will never need more than 15 megabits per second between point A and point B. As a typical service provider, I could tell you that I will use an MPLS, bandwidth-on-demand, optical switching model to guarantee 15 megabits per second from point A to B. However, my theory is that today's GigE equipment costs and dark fiber can enable you to tell the customer that despite where 0.2% of the time you might want to peak at more than 15 megabits per seconds, you'd rather sell him GigE services at a 20% higher price than a competitor would provide him with his exact 15 megabits per second. I take the risk, as a service provider, that it's indeed worthwhile for me to service that customer with 985 megabits per second capacity than he does really require on the basis that he may peak at more than a 15 megabits per second less than 0.2% of the time.

COOK Report: And the cost of servicing that extra 985 megabits per second capacity would not cost a lot more than 0.2% above what a competitor would sell this 15 megabits per second service for?

Menard: Exactly, hopefully at no more than a 20% premium. What I hope that the customer will understand is that that this 20% premium enables him to empower adaptive applications with the 985 megabits of additional capacity in addition to providing him with a serious buffer for a few months of growth in his bandwidth demand forecast. My ability to make the customer understand will be entirely proportional to the adaptive applications that I can empower on his network as an eCommerce Application Service Provider, hence why my next task at VDN is to build the ASP venture.

Instead of having designed a network for applications, you can design your applications to be adaptive to a totally over-provisioned network. Imagine that you have access to all of that excess capacity all of the time, without any impediment in the way of using this capacity such bandwidth sold at a premium above committed access rates pricing.

COOK Report: And the advantages of doing this? I mean, how might you go about it?

Menard: Because of extensive fiber capacity and low cost structures, VDN can apply inexpensive equipment and generate ten gigs of bandwidth for the amount that competitors would have to spend to generate a single gigabit.

COOK Report: But it sounds like you are taking a situation of VDN and generalizing it to the entire Internet. You want to apply your model to metropolitan areas, which because of the fiber builds of the last five years tend to have lots of fiber. Currently cheap gigabit Ethernet allows you to provide a gigabit of service for what a carrier/SONET based service would charge for 10 megabits. But in just how many cities is this business model applicable? Probably in a lot more than one would have thought a year ago. One wonders whether it is feasible in most cities? What do you do?

Menard: Simple, build gigabit and eventually 10 gigabit networks where fiber or right of way's are accessible. It would not be practical to apply the 10GigE/NO MPLS model in an environment, where putting in a lot of strands of fiber is too expensive. This is why the gigabit Ethernet model will not be ubiquitous and there is likely to be a lot of MPLS. The problem is that MPLS will not help the situation where the backbone has to be 7 orders of magnitude faster than the PSTN because the access networks will become 5 orders of magnitude faster than the PSTN.

Meanwhile to go back to what we are doing in Montreal, when you have access to a network like VDN's you can convince your customers that it's worthwhile to have extremely adaptive applications on the network. In such an environment you know what, the worst case situation is the minimum bandwidth necessary for all applications to function properly. You also find that when there is extra available bandwidth, then you can do all sorts of neat things.

But the problem is that the existing set of real-time applications is not adaptive and, even if there were extra bandwidth, wouldn't use that extra bandwidth. I think telephony is a very good example. If there was extra bandwidth, nobody would run high fidelity audio, because everybody's talking with crummy POTS or IP phones which are built up to pump out no more than traditional PSTN "toll-quality" sound.

I am saying that if we begin to get successful implementations of affordable multi-gigabit IP infrastructure in many cities, it will create an environment which will serve as a source of motivation for people to develop what I am calling adaptive applications. This

environment will make it possible for people to experiment with many new uses of audio and video interfaces for the web, for Internet telephony and for many other applications. These new applications will need to be written in ways that their feature sets can be throttled back when network congestion increases. These new applications will be able to adapt to network conditions.

For example in a high bandwidth adaptive environment, you might have a hierarchical codecs. If additional bandwidth were available, your audio applications would make use of the additional bandwidth chunks by increasing their fidelity and allow you to listen to the higher bandwidth multicast sessions rather than the lower bandwidth ones.

COOK Report: And if the network starts to crunch, they can throttle back down.

Menard: They can throttle down. And this is where I think my logic is very difficult to understand for a Bellhead. If I were a Bellhead, I couldn't understand how a customer could deploy effectively adaptive applications. Especially adaptive multimedia applications. Why? Because I would assume a constant scarcity of bandwidth. But the Bellhead ignores the TCP protocol's ability to adapt to network conditions by backing off when congestion begins to appear, and this Bellhead certainly does not regard adaptive multimedia applications as being anywhere necessary or commercially viable.

COOK Report: Can your applications adapt by just an increase in their numbers? Instead of ten people doing some sort of video link between point A and point B, having twenty people a month from now doing that link and fifty people two months from now? Or they could also adapt by having the same ten people next month doing it with a higher level of fidelity or more colors?

Menard: Exactly. You always move upwards in an Internet environment. In the Bellhead world you simply have your baseline voice service and for as long as you envision that the normality is the baseline service, you can't envision anything else. But the thing is, on an Internet, when congestion hits, you throttle down to what the Bellheads consider being the baseline service. And with multicast technology, you only need to send the traffic downward once.

COOK Report: Okay, but let's take a hypothetical area, where instead of getting fancier applications or more bandwidth intensive applications, the person envisions just more people doing the same kind of applications from point A to point B.

Menard: I would answer by making another reference to my hypothesis that by providing GigE service for those that may need

ten times less bandwidth, a very large headroom is being provided to permit more users to do the exact same thing. It is both the combination of adaptive applications and extensive additional capacity that end-up tilting the dominos in favor of the value proposition that I am describing.

COOK Report: Are there some working groups developing some new, adaptive protocols now?

Menard: These protocols are already in existence. It is just a matter of using all the capabilities of the Internet Real-Time Protocol (RTP).

COOK Report: So, in other words, it's a matter of people taking some audio codecs and expanding their capabilities.

Menard: Yes. But voice is not even the killer paradigm shift to illustrate the necessity for adaptive applications. The killer application is possibly the multicasting of high definition TV (HDTV) GigE networks, versus, for example, MPEG 2 and then MPEG 1 video quality, where Internet TV sets will throttle down to potentially something that is similar to an NTSC quality in situations of congestion.

COOK Report: And in getting them able to throttle down does not involve anything all that fancy? I mean, what does it involve? Coding, some means of going from one codec to another? Given the state of TCP/IP congestion? When it hits a situation where packet queues are two packets long or they don't exist or some other factor?

Menard: Yes. For example, you can set your network, whenever there is more than 75% link utilization, you have to unsubscribe to HDTV multicast groups and live with no more quality than the traditional TV environment.

Another thing is, by that time, when you start thinking about this, you should not be thinking about real time streaming of everything. When you're dealing with GigE networks, you start to understand very quickly that video programming built of play lists of files may be more efficient than streaming because it can use maximum packet sizes all the time.

COOK Report: So this really is movies on demand.

Menard: I would call it downloading a movie to your hard drive. You would have a set-top box with a hard drive and when you say I want to watch this movie the network would start to download it and would give you an estimated delivery time. From the point of view of network congestion, what's beautiful about downloads of movie files is

they are a much more adaptive as an application than they would be if you wanted to view them in real time.

COOK Report: Then you're saying, instead of having an application going in real time from point A to point B, the idea is that, if it is a video application and is transmitted all the way to point B, you can, before the person at point B begins to look at it, use it, observe it, tweak it?

Menard: Yes, or give you at least a good fifteen minutes head start in doing so.

COOK Report: Like a fifteen-minute buffer.

Menard: Yes. So that even if you need to throttle down because you've suddenly hit congestion because too many people are doing true real time applications, then you don't congest the network with video downloads.

Two Schools of Thought

There's two schools of thought. You either build a network for adaptive applications and then you as much as possible design applications to be adaptive, or then you make an assumption that you're dealing with a network where it makes no sense to build adaptive applications. Say you're trying to run a mobile phone network on IP. Ninety-nine percent of the traffic will be real time. There is no hope at all that you'll ever get adaptive applications on a cell phone. So what would be the point of throwing all of this onto the Internet?

If you are a network operator that's not entirely an ISP, you will probably not care about integrating that traffic onto your IP network. It is precisely for these people that an MPLS network becomes interesting. You put that network on a different label and then let that co-exist with your existing labeled traffic. However, by using MPLS, you can leverage your IP routing expertise and do not need a separate management network. But my suspicion is that the networks for carrying these types of non-adaptive traffic are already built. While these networks are growing, they are doing so at nowhere near the growth rate of the Internet.

The point is that I am not too worried about an extensive migration to MPLS. Instead I expect the building out of much faster Internet backbones with massive bandwidth in the core that will once again make the complexity of managing virtual circuits something too complex to be reflected in a pricing scheme that can involve the end-users. My vision is to build a bigger, faster Internet that will be future-proof. By this I mean that I will do everything to ensure that the applications which will be running on

that future Internet will be extremely adaptive.

COOK Report: I can understand this at one level. For example, I've been using PageMaker on my newsletter and I got Adobe In-Design and experimented with that and I found the size of the file grew almost by an order of magnitude. Instead of 200 kilobytes, it was two megabytes. So in one sense, over a period of time, a lot of applications will tend to become more bandwidth-hungry, at least what we're used to and what we expect.

Menard: Yes.

COOK Report: And in this sense, if you have your networks that can easily adapt to using the super-duper quad definition whatever bandwidth hungry application, when things get a little crowded, you go to only thousands of colors instead of millions and presumably nobody gets too unhappy, is that part of what you're saying?

Menard: And the the winners will be service providers who will be able to recommend and enforce newer kind of application policies instead of operating a box which penalizes everybody when applications don't fall-back to lower bitrate codecs. The new network elements will provide direct feedback to end-users which will be forced to switch back to a less bandwidth hungry codec when instructed to do so, or face the possibility of being shut down.

Here is my point. As a network operator, if I've got congestion then I start watching people. I don't have to care when there's no congestion. If I have enough bandwidth, I will not have to worry about watching for congestion on a constant basis. With adaptive applications, when I find the offending application, I can offer it to accept my recommendation to run at lower bandwidth or else I will have to take punitive measures. I find this process friendlier to Internet applications than resorting to the use of application-aware layer 7 switches, which may render the user experience very unpredictable across several applications under different conditions. It may fly in the corporate environment, but for residential services, it is nearly impossible to explain end-users to expect very different experiences across different applications. Nortel (Alteon Web Systems), Top Layer and Foundry Networks amongst others make such equipment.

COOK Report: Might there be an override where, if they don't obey after a certain period of time, they get a warning, you're being timed out in ten seconds?

Menard: As a user, I would prefer this anytime over having a Layer 7 switch start prioritizing access to certain ports. By mandat-

ing user intervention, I do not have to discriminate against one application instead of another. I discriminate against the IP activity from one particular user behind one particular IP address.

COOK Report: But let me back up and say that I'm still, maybe your Neanderthal MIS person and I say, well, having quad TV definition of the chairman's meeting in real time between Montreal and Miami, Florida, doesn't seem that exciting to me. What I want to do is to have some growth room if, over the next month, they double the number of people sending ordinary IP voice 8-kilobyte codec sessions from Montreal to Miami, Florida, I want to be able to cope with that. That, is the kind of network growth I'm expecting.

Menard: I would answer him that whereas he could buy a T1 today, I would provide him with GigE for not much more. You give him a network that's more than 600 times faster. He will not have to worry about such growth on the access network. However, the Internet backbone bandwidth will be forced to cope with this massive increase in access bandwidth. This is where the whole argument lies now.

COOK Report: But what confused me is I thought you were telling me initially that he could only use that extra bandwidth capacity by having super-duper adaptive bandwidth applications rather than more employees doing the ordinary things. What if he says I got to be ready for more employees doing plain vanilla stuff?

Menard: That's a very simple problem for me, because that's not the big problem. My big problem is, what if instead of having twice as many employees doing the same thing, you have 40 employees all of whom understand that now all this bandwidth is available and start running really bandwidth-intensive, non-adaptive applications. Because to me, the problem is not so much the network growth as it is the ability of certain new customers or end users in your organization to start sucking all of the available bandwidth.

COOK Report: Are you are saying this to contrast with typical pricing schemes from the tier one Internet backbone providers of today? That is, if I need 45 megabits per second from point A to point B and I have that pretty well worked out, you would say to me, well, your business is growing, you're going to use new applications, so you're obviously going to outgrow this 45 megabits per second. Then you're going to say to me, look, Cook, I've got a real good deal for you, I can give you ten gigabits for a price that's only 30% more than you've been paying for the 45 megabits. Think how nice it will be that you don't have to waste time

every week calling up for a new leased line that you suddenly need because of your increased bandwidth consumption.

The Function of the MIS Manager in the 10 Gigabit Ethernet World

Menard: Exactly. And then instead of me charging you for bandwidth on demand, you can start using adaptive applications that can make better use of that available bandwidth.

COOK Report: I'm still fixating on how that extra bandwidth is used. I can start having employee growth if I want or some employees can become bandwidth hogs and start running multiple applications, etc.

And I understand the importance of what you're saying about adaptive applications and why people want to become more cognizant of this as a possibility. But at the same time, in this model that we're talking about, there is nothing that prevents the customer from beginning to use up the surplus bandwidth with just growth in number of people that he's got doing things on the network.

Menard: True. However, I stress that the MIS manager need to enforce organizational policies for use of certain Internet applications. In some sense, MIS managers have to worry about the applications that are loaded on computers the same way that they have to watch for software licenses.

What I would not expect to happen is a customer starting to complain to me as the service provider, that the connection is not fast enough because his MIS manager is not aware that 10% of the users on their network are causing 90% of the congestion on a GigE link. If this happens, these people will be easy to find. So the faster the network, the more I would expect a network manager to be cognizant of the fact that, if the network starts to have congestion, it's because he cannot enforce company policies.

COOK Report: Okay, and if that network starts to have congestion, I, as the network manager, have the choice of going to and buying even more bandwidth, right? That's one choice?

Menard: Yes, and again, you'd be purchasing it in increments of Gigabits per second rather than megabits per second, which make a big difference in the frequency of your procurement for additional bandwidth. More importantly, you may also have the opportunity to enforce certain organizational policies by watching what applications are installed on users computers, by way of normal software inventory processes. My point

is that if a couple of users are going to cause 90% of the congestion on a 10 GigE link, it will be because the MIS manager will have allowed them to load application which are not in compliance with corporate acceptable use policies. Humans should be enforcing policies on other humans, rather than asking machines to discriminate against a certain set of applications.

COOK Report: But if for any reason the MIS manager's policies as to individual use didn't work, another way that the manager could do this is to replace a layer 2 switch out there with a switch that's Layer 7 aware.

Menard: Those boxes are proxies for a MIS manager's inability to supervise the type of applications that are allowed to run on the company's networks, particularly in environments where network bandwidth is scarce

COOK Report: Understood, but as long as they exist, let's tick them off real quick.

Menard: What I am saying is that in a GigE context, those network managers have the opportunity to manage their networks through traditional software inventories on user's desktops by removing applications which are not certified to respect organizational acceptable use policies. The use of Layer-7 switches is a poor technical alternative to solving the problem at the source, and is difficult to conceive of in residential deployment contexts.

COOK Report: So while you do not recommend this option, if a company were using it, the box would be out there on the network and you might have it set such if Cook's application starts hitting the network too hard, the network will shut it off.

Menard: The objective is to control applications (and users behind applications) that behave in non-adaptive ways. These boxes are application-aware, but never send feedback to applications. Their means of control is not enforcing control over applications. Their means of control is to start killing packets on the wire, so that non-adaptive applications do less damage to the network. But it will never, ever replace a network monitoring box that shows that certain applications are not behaving properly, and send feedback to the applications telling end-users to behave properly in some automated manner, penalizing offending users rather than certain applications.

COOK Report: Let me summarize. You involve the end user by saying you're using too much bandwidth, either throttle down or your connection will disappear in one minute?

Menard: Exactly. Instead of controlling application in a master-slave fashion, network ad-

ministrators can trust applications on the basis that they are known to react to network congestion warnings and will shut down automatically if the user does not react appropriately.

COOK Report: And the other way that layer-7 switches do that is not by providing that feedback, but rather by just dropping packets.

Menard: Exactly. It never gives back feedback to the end user.

COOK Report: Your point is if you're going to use something like a layer-7 switch, the use of which doesn't give feedback to the end user, you presumably have it in place only because the manager is failing to give users the necessary feedback in the first place.

Menard: Exactly.

COOK Report: But just to cover all the possibilities, could you also do this with a media gateway controller?

Menard: Instead of having applications, which are controlled by end users, you could have applications, which are controlled by network elements. And a typical way of doing that in a voice-over-IP context is by using a master-slave type of protocol, such as H.248/MEGACO (the MEDIA Gateway Control Protocol designed jointly by people participating in the ITU-T and the IETF). This way, the by having the telephone being controlled by the network, you can be sure that it won't run applications that the network won't be aware of.

Either you provide a feedback path to the end user and the end user controls himself or you end up controlling the whole terminal on behalf of the end-user. But suppose I am an end user and I load on my PC this totally cool, latest generation, quad HDTV-quality streaming television application which is also non-adaptive. If I load this against my manager's wishes and I start disregarding my network use policy box that's telling me, hey, there's congestion here, slow down, buddy, then I'm going to get my hind end kicked.

COOK Report: And it's going to be kicked by the media gateway controller?

Menard: No, it cannot. That's why I find that resorting to control of terminals rather than enforcing proper software inventory management is somewhat of a kludge. Media gateway controllers have to control the whole terminal and prevent end-users from loading applications that are capable of bypassing media gateway controllers or gatekeepers. To give a practical example, on my PC, the MIS person in my company may configure Microsoft Netmeeting to only make calls when receiving permission from an H.323 Gatekeeper, but if nothing is stopping me from loading this latest SIP-based Internet Telephony application that totally disregard the existence of an H.323 Gatekeeper,

then he can certainly advocate for the use of a Layer-7 switch. However, if I were forced to use an IP phone on which I could not have the right to load additional applications, and then there would not be a problem, would it?

COOK Report: But we are still going to see something like Victoria's Secret having their fashion show on the Web or some big rock group having a webcast of a concert somewhere. Something that would spike Net bandwidth tremendously would be real time webcasts of this type of thing. In such a situation given number of people in the office might want to open a window on their screens and view the web cast at the same time that they have the more pedestrian things running?

Menard: Well, yes, the user network is going to be marked by these unique events. As an MIS person, I would not dare to go to my boss and ask him to OK the budget to purchase a layer-7 box on the basis that it is needed to balance traffic when the employees of the company are watching Victoria's Secret webcasts. It boils down again to the poor control that MIS persons have over software inventories inside user's PCs. You can think that the content adviser systems inside operating systems could be slightly improved to collaborate with software inventory management infrastructures. I find it to be very intriguing that the Microsoft System Management Server (SMS) has no interaction with the content advisor system built into Windows at this time.

What I refuse to accept is to see people not understand that the future inevitably points towards multimedia applications that can and should be tailored to adapt to network conditions. Such applications will be able to auto-configure themselves to run according to the amount of bandwidth that is allocated to them at any one time by organizational policies.

COOK Report: What I was having trouble seeing is reasons why people who are responsible for making these applications available should have them to be adaptive. But if we could agree that webcasts of something like the famous Victoria's Secret fashion show and/or webcasts of the surviving Beatles getting back together are important, your idea would be that there should be multiple means of tuning these in and one of them should take only so many kilobits per second, so many hundreds of kilobits per second or so many megabits per second.

Menard: Yes, I agree. Again, my main motivation in what I am doing now is to understand the implications of adaptive versus non-adaptive applications. Sorry for coming back to that, but it's truly my opinion at this time.

If I sell you a GigE or a ten GigE service, whereas I do not ever expect you to use anything more than fifteen megabits per second, then you've got access to all of that available bandwidth. What I then reject is the business

model of the carrier that says that he will sell you flat rate bandwidth and allow you to exceed that slightly without further charge but that if then, if you exceed beyond some defined threshold, he will be entitled to severely charge you, because he has taken the decision to over provision at his own expense. That argument is simply bogus, as I find that it does not cost tremendously more to over provision 15 megabits per second with GigE service to begin with if the fiber capacity is there

COOK Report: Well, this is the UUNET pricing model that you're getting at, that you're charging by usage at the 95th percentile.

Menard: Yes, instead of the customer being given the privilege of running adaptive applications, their service provider is essentially saying, Sorry, Mr. MIS Manager, you cannot afford to run adaptive applications. And, by the way, if you ever peak past that one megabit per second, the bandwidth that I have over-provisioned you with will be billed you on demand with a hefty premium. So its essentially, the upstream service legacy carrier-based Internet service providers (the Tier-ones) which are forcing me to sell access to the global Internet this way. Once new backbones become available, it is my hope these business practices will be done-away with. For those like us that have chosen to focus on the access network, the opportunity to develop such artificial pricing schemes needs to be carefully weighed-in against the lost opportunities to rapidly deploy new adaptive applications.

COOK Report: And all of a sudden, you'll go from in effect being flat-rate billed, to not only being flat-rate billed, but for your peak usage, you're being billed according to dynamic bandwidth requests which is going to screw your budget even more.

Menard: Possibly and you're paying a premium for that flexibility. I find this to be very disturbing to see carriers attempting to convince MIS managers that buying into MPLS is a wise decision for the benefit of customers. I find this to be nonsense given the economics behind GigE metropolitan optical networks that VDN is currently engineering in Montreal.

COOK Report: So one of the major conclusions of this discussion is that Montreal is only one example, and that there are going to be more and more cases where dark fiber and gigabit Ethernet will provide good opportunities for replicating the VDN business model that we've been talking about.

Menard: Yes, exactly.

Afterword - We have text of a long second half of the **Menard** interview on the implications of what has just been discussed for the global backbones of the Internet. Having withheld it to make room for the Howard Frazier interview, we may publish it in a future issue.

10 Gigabit Ethernet Draft Almost Out

Gig & 10 Gig E Enable SONET Free IP Data Nets

Availability of Cheap Transport Technology for Both MAN and Wan Enables Enterprises to Drop Carrier Managed Services

Editor's Note: Howard **Frazier** holds the title of Distinguished Engineer at Cisco Systems, Inc. within the Gigabit Technology Group, where he is responsible for product functional and architectural definition. He is one of the inventors of 10 Gigabit Ethernet, and a key technical contributor to the IEEE 10 Gigabit Ethernet standards activity.

He was the chairman of the IEEE 802.3z Gigabit Task Force, which wrote the Gigabit Ethernet standard. Previously, he was the chairman of the IEEE 802.3u 100BASE-T Task Force, which developed the standard for Fast Ethernet. He holds patents on both Gigabit Ethernet and Fast Ethernet technology. He is a member of the IEEE Standards Board, and the IEEE 802 LAN/MAN Standards Committee. Prior to joining Cisco, he was employed by Sun Microsystems, Inc. He graduated from Carnegie-Mellon University with a BSEE in 1983. We interviewed him on July 28, 2000

COOK Report: Since we interviewed Dan Dove of Hewlett Packard last December on gigabit and ten gigabit Ethernet, would you start off by talking about any significant developments since then including developments in the standards process?

Frazier: The standards process in the IEEE is long and rigorous and within the 802.3 working group which does Ethernet it is especially rigorous. Ethernet is and has been way far and away the most popular local area networking standard in the world. There is a lot of interest and very large amounts of money involved in the process. Because the standard is so very widely used, we have to be careful. It is very tempting to listen to any new proposal that comes along and say: "yeah we can do that, because someone thinks there is a market there."

COOK Report: But you have to pay attention to issues of backward compatibility?

Frazier: Yes. And the 802.3 standard in paper form is a good two inches thick. When people say they have great ideas, one of the things we insist that they demonstrate is that their ideas have broad market potential. Our informal position is that you better be able to convince us that you can sell millions of these things. We try very hard to protect its affordability, simplicity and ease of use.

Consequently when we do a standard, we do it very very carefully. At the same time we do them very quickly. The fast Ethernet standard (ten to one hundred megabits) from start to completion, took 2.5 years. The gigabit Ethernet standard from first mention (November 1995) to completion (June 1998) again took about 2.5 years. In March of 1999 we started a ten gigabit Ethernet project. It was approved by the IEEE standards board in December 1999. At that stage we had defined the objectives of the project. We had a one page summary of what we were going to do. We decided to support both a LAN operational mode for 10 gigabit Ethernet which would operate at 10.00 gigabits per second and WAN operating mode running at a data rate compatible with the SONET OC-192 rate which is 9.29419 gigabits per second.

Key Issues in the Standards Process - Jumbo Frames

COOK Report: One assertion which one sometimes sees is that to be useful with IP at these higher speeds, you must have Jumbo frames that would carry more packets.

Frazier: We decided that we would not support jumbo frames. I call them mutant frames because I am vehemently opposed to their use. The idea of supporting large frames has come up every two years in the 802.3 committee. It comes up because it makes people's benchmarks run faster. If you run a point-to-point test between two computers, the bigger the frame you use, the faster your benchmark will run.

COOK Report: The more data you can transfer in a given amount of time?

Frazier: No. It doesn't do anything for efficiency on the wire. It may increase it by a percent. Maybe two percent.

Here's the issue. If you can manipulate a page of data in computer memory at a single time, then you can do some memory management techniques which are quite efficient. In particular, when you receive a page of data from the network, where a page in a computer is normally 4,000 bytes, or 8,000 bytes, you can process it very efficiently.

This comes down to the memory management techniques that people use in computers. You may use jumbo frames that are at least the size of a computer memory page plus room for TCP/IP headers. Therefore, if you want to accommodate a four k page, plus your TCP/IP headers, and you have a packet (frame) size that is at least five k, you can transmit an entire memory page plus the headers necessary to get it to your receiving host. When the receiving end gets the jumbo frame it strips the headers and blips the page up into the user space without having to do a copy to memory or any re-assembly.

COOK Report: One of the reasons for wanting to do this is that network speeds are now faster than CPU speeds?

Frazier: Right. Now back in the early days of Ethernet, when it was a research project at Xerox, they didn't have a machine anywhere in Xerox that could saturate a 10 megabit per second Ethernet. It was the mid 1980s before a single machine could do this. A Sun Sparc station couldn't saturate a ten megabit link until about 1989 or 1990. When FDDI first came out, it was a struggle for us to saturate that. However when fast Ethernet first came out many commercial computers could saturate it. But in the global scheme of things you could say that fast Ethernet came along rather late - some 20 years after Ethernet. Later when gigabit Ethernet was implemented, no one could even begin to saturate it.

And when the hardware makers found that their big super servers couldn't saturate a gigabit Ethernet, the only way to make their benchmarks look impressive was to create jumbo frames. Then they could do memory to memory tests. But these are totally artificial benchmarks. Nevertheless when they could get close to a gigabit per second from a large server, they were happy. This was the cause of the drumbeat for jumbo frames.

COOK Report: So you would say that it is a way of giving a misleading impression that a machine can absorb a gigabit per second from an Ethernet in any meaningful way?

Frazier: That's right. What has happened over the past couple of years is that computers have been getting faster. Protocol implementations have been getting better. People have figured out how to build smarter net-

work interface cards. They have figured out smarter techniques for doing direct memory access that make the memory management and device driver implementations much more efficient. Consequently a high end Sun Microsystems server can easily saturate a gigabit Ethernet today. Moreover it can do this without using jumbo frames. Therefore there is no need for jumbo frames.

Here is where they breakdown. I said they were only good for benchmarks. The reason for this is that the problem is only in the receive direction. In the transmit direction the small frames don't hurt you because you can take a big page of memory and very readily and simply chop it up into packets. Computers can out put a gigabit very easily. It is the sucking it in, the input or receive process where they slow down.

Think of a file server in a client server relationship. A file server typically transmits information to clients. But when you want to transfer a file to a file server, first of all you find that the data is coming from a desktop somewhere in the network. However, unless you get all of the desktops in the world to support jumbo frames, it doesn't really do any good. Of the 200 million desktops in the world that are on Ethernet zero percent them are currently physically able to support jumbo frames. With clients still operating at ten or 100 megabits per second, nothing can support the jumbo frames.

COOK Report: But what about Ethernet done over a wide area network?

Packets per Second as Key Measurement

Frazier: Do routers originate data? No. They pass on data that they get from a host computer somewhere else. It is true that one of the ways in which people are attacking this jumbo frame issue is that they are going to an IETF working group that is specifying the IS-IS routing protocol. They are having that group standardize jumbo frames because you can argue that it's good to send jumbo frames around to exchange routing protocols. There has been an RFC to standardize jumbo frames. The RFC was by Alton Web Systems one of the proponents of jumbo frames and intended to try to get jumbo frames for Ethernet specified within IETF. This is never going to fly because the IETF does not control hardware specs.

You might say that as you are passing gigabit or even ten gigabit Ethernet between Internet core routers or switches jumbo frames might help you. But the other point here is that switches have absolutely no problems keeping up the packet rate for gigabit Ethernet or even ten gigabit Ethernet when

you are doing minimum size packets. You see what counts for a switch is not how many bytes per second it can ship but how many packets per second it can forward. Switches must be designed to be able to forward the smallest possible packets at wire rates. That is the design limitation on switches. Sending big packets is easier for switches because they are never ever bottlenecked at a byte per second rate.

COOK Report: Does the same thing hold true for routers?

Frazier: It depends on the router implementation. Software based routers like the Cisco 7500 does software with the same limitations that a typical file server suffers. Big packets would probably make such a router ship more bytes, but it turns out that this router in particular is limited by its memory bandwidth. It can't forward Ethernet any faster than at about 300 megabits per second.

The large routers now bifurcate the process doing the routing table lookups in silicon and doing data forwarding while the result is sitting in memory. You may rely on software for strange headers, but typically in the bigger routers you are doing all this as a layer three switch.

COOK Report: So the operation of the largest Juniper and Cisco routers are not dramatically different from switches?

Frazier: That is right. The point is that switches always have to be designed around the number of packets per second that they can process. In that environment jumbo frames won't make one whit of difference. For gigabit Ethernet, if the switch is going to work at wire rate, it must be designed to forward 1.488 million small packets per second. Make the packets bigger and the switch forwards 81,000 maximum size Ethernet frames. Make them jumbo frames and cut the number to ten thousand. Is it somehow easier on the switch!/? No. It doesn't make any difference to the switch. Ramming larger frames through the switch does not increase its throughput at all.

Now if you introduced jumbo frames on high end servers that were attached to the network at say gigabit or ten gigabits, you'd have your servers talking big packets., remember that you have 250 million clients that will not support these frames. Their silicon and their drives cannot cope. How are going to upgrade 250 million NIC cards? So if the servers talk big packets and 250 million clients talk little packets, then something in between would have to fragment the big packets into little packets and in the reverse direction going from client back to server would have to re-assemble them.

COOK Report: In such a case you are talking the expense or more equipment which itself can fail as well as the process of fragmentation and re-assembly vitiating any efficiency gained from large packets?

Frazier: Yes. Now we have also had the lessons gained from experience with this. FDDI supported 4500 byte packets which are three times the size of the largest Ethernet frame. In the early 1990s lots of people were hooking their servers up to FDDI with desktop clients still on Ethernet. People built boxes that translated from FDDI to Ethernet. A lot of those boxes attempted to do fragmentation. To convert the FDDI frames into little Ethernet frames. And packing little Ethernet frames into a big FDDI frame. They had so much trouble getting this to work that by default every FDDI server just turned their MTU (Message Transfer Unit) size down to Ethernet size. And when people installed LANE, ATM with LAN Emulation, they did the same thing. They put ATM in their servers which could support big frames. But people found they had to do the same thing — go in with software and turn the ATM frame size down to Ethernet size.

So the morale of the story is that jumbo frames are a bad idea and we don't need them at gigabit or ten gigabit Ethernet. What we do need are better protocol implementations, better network interfaces cards, and memory management implementations. All these things are being worked on.

Standards Process

COOK Report: To pick up again with the standards process?

Frazier: Getting back to where we have been for the last six months. We have spent the first half of the year 2000 in hearing detailed technical proposals designed to achieve the objectives of the standard. We have had hundreds of technical proposals. There are about 250 who show up for each meeting. At any given meeting perhaps a third - about 80 people - are new and have never attended a meeting before.

What we have underway then is a continuous updating, coalescing and refinement of proposals. We try to take the best ideas from multiple proposals and meld them together often by asking people from two different proposals to work together. Our end goal is to be completed in March of 2002. We expect to have a first draft of the standard in September 2000.

We have just concluded a key meeting from July 10 to 14th. At this meeting we formally adopted a baseline set of nine proposals which purport to tell us how we will reach

each of our objectives. Between now and September these base line proposals are to be converted into a first draft. We have a team of about a dozen editors who will take this slide ware (Power Point presentations) and convert it into the language of a draft standard.

However there is one particular area where we have been unable to select a baseline proposal. We want to be able to support multi-mode fiber. This is the fiber that is most commonly used inside buildings for building backbones within an enterprise. This fiber itself is actually more expensive than single mode. But the connecting hardware for it is generally cheaper than for single mode fiber.

This was true, because for slower speeds, 10 megabit, 100 megabit and even up to 622 megabit ATM, you could use light emitting diodes (LEDs) to drive the multi-mode fiber optic cable. The problem is that once you get beyond OC-12, LEDs cannot be modulated fast enough. You need to use a laser. You can use a laser to drive multi-mode fiber, but it turns out that this same laser can also be used to drive single mode fiber.

Now we have to establish what we call economic feasibility. We use a guideline that we want to provide ten times increase in performance or no more than three times increase in cost. You get ten times the bang for three times the buck. What we find is that over succeeding generations of the technology, the cost of the new generation will rapidly decline to somewhere near the cost of the old generation. Look at it this way. It doesn't cost any more to buy a ten one hundred adapter or switch port now than it does a ten megabit only port. They are about the same price. Right now gigabit Ethernet components cost what fast Ethernet components cost four years ago.

The price of fast components has gone down even more. The price differential from gigabit to fast Ethernet is something like two to one now. Ten gigabit is quite a challenge because at this stage in the game the components are quite expensive. But this is the promise of Ethernet. We will find a way to get the volumes up, get the simplicity into the system, get the complexity out and get the cost down.

COOK Report: So the debate in this area is over precisely these economic issues?

Economic and Multimode Fiber Issues

Frazier: Yes. The bang for the buck. Ethernet provides the cheapest cost per bit of any transport technology. In terms of

dollars per megabit per second none of the other technologies such as ATM or FDDI comes close.

COOK Report: Perhaps the uncertainty comes from your ability to project and estimate which way of achieving your design goals will be the least expensive?

Frazier: Yes. But some of the other variables that are thrown are simplicity and ease of use. We try to make any set of solutions that we pick satisfy the broadest set of applications possible. We don't want to have to provide too many options to users. For example there are at last five or six competing proposals for how to support this multi mode fiber optic cable in buildings and people have a variety of ways to do it.

There are some people who want us to adopt all of those ways and include five or six different physical layer, different connector specifications, different port types, different optical parameters in the standard. None of these are interoperable by the way. But if you put them all in the standard and vendor A choose to implement one and vendor B another, you wouldn't be able to hook vendor A to vendor B. It would also be very unclear when it was best to use each one.

The over riding message from system vendors is that we want the a set of the fewest options possible. We want to make it very clear cut as to what the standard really is. It is not a standard if there are five or six choices. It would be a menu. A standard says here is the best way to do it.

COOK Report: And someone who advocates the menu approach may be doing so because they have some specialized situation for a specialized market?

Frazier: They may have a technology. One of the things that has happened is that a lot of the drive for having a multiplicity of solutions has come from manufacturers of fiber optic components - both cables, and connectors and optical transceivers. Why do they want so much differentiation? Because they all wrote business plans based on those things.

COOK Report: You have a lot of people developing a lot of things in parallel. Here is another industry developing its components at the same times as ours. This is part of the complexity.

Frazier: Yep. And where we are at right now is that because we have this multiplicity of choices for multi-mode fiber we haven't been able to make a decision on how we are going to support multi-mode fiber. As far as single mode fiber is concerned, we have two decisions. One phase will support single model fiber cabling distances of up to ten

kilometers on single mode fiber and that is intended primarily for campus applications and for connection into DWDM transponders. The same spec will be good for hooking a switch or a router up to a DWDM box. The other thing that we selected was a specification for going long distances on single mode fiber. These are distances of greater than or equal to 50 kilometers.

There are many other areas in terms of how the logic works, how data is encoded on the medium, and things like how packets are formatted where everything is agreed upon and settled. The base line proposals here are being written into the draft so that people may review and comment. I am hopeful that the decisions on multi-mode fiber will be completed by November of this year.

COOK Report: What percentage of currently deployed fiber is multi-mode?

Frazier: It depends on what application space you are in. If you are looking at a building backbone there is only a small percentage of single mode. Five or even ten to one multi mode to single mode. In a campus environment it is about evenly split between multi mode and single mode. When distances extend over more than about two kilometers you can't send data on multi mode fiber without the amount of attenuation and dispersion in the fiber just killing you.

Ten Gigabit Ethernet Products are Here

We have left building application space wide open. But for companies like Cisco this is the most critical space. Cisco has publicly announced ten gigabit Ethernet products. At the Las Vegas Inter Op show in may of this year we conducted a demonstration of a card or "blade" for a Catalyst 6500 switch. This is an enterprise class Ethernet switch. The blade that we demonstrated in that switch had two ten gigabit Ethernet ports on it. Moreover the ports can support various different types of optical interfaces. Because of the uncertainty in what the optical interface will be in the final standard, we made the product flexible enough so that we can replace the optical components when the standard is finished.

COOK Report: What then does this switch with the two ten gig Ethernet ports enable the Enterprise to do that it could not do before?

Frazier: What we are seeing is a trend in which the speed to people's desktops and their servers is increasing. Over the last year most desk tops have become connected to the LAN at 100 megabit per second speeds rather than ten megabit. We are even begin-

ning to see the deployment of gigabit Ethernet to the desk top using unshielded, twisted pair, category five cable.

In general what we find is that, if the desk tops go at speed x, the servers need to run at 10x and the backbones at 100x. The backbone after all has to aggregate traffic between many desktops and many servers. People are moving away from distributing their servers at the work group level. Rather than departmental servers with highly distributed data not under any central control and given the increasing importance of the data on those servers, managers decided to centralize the servers within an enterprise in so called glass houses almost like a return to the mainframe computer room. Under such an architecture you would have to transit the backbone to get from the desktop to the servers.

COOK Report: These developments seem to parallel those of remote application service providers where much remote storage may not be in the next room but hundreds of miles away.

Frazier: True and when you wonder why is Cisco offering ten gigabit Ethernet you will find that the answer is that the speed of the end user servers is getting faster and the speed of the backbone has to stay ahead of the speed at the desk top. With increased bandwidth at the edges the demand for increased backbone bandwidth is growing. The demand for ten gigabit Ethernet is primarily for use in backbones connecting enterprise switches to each other. Now such switches can be sitting in the same rack, or in the same building or across the campus or across the continent from each other. It then goes from switches to routers and the next stage will be to send it from a switch to a router or server - although it will probably be a couple of years before there are any servers that can do serious things with ten gigabit Ethernet.

The bottom line is that we are seeing customer interest in and demand for ten gigabit Ethernet. Cisco has committed to ten gigabit Ethernet products that will ship before the end of this year. They are not IEEE standard compliant because, as of this point in time, there is no IEEE standard. Cisco will offer upgrades to standards compliant hardware. We are making the physical layer interface and the optical modules in a modular format which will allow for easy replacement. Right now the customers will buy these because they have a need for ten gigabit interfaces. While using it in a pre standards format, they will benefit and achieve some cost savings somewhere else in their infrastructure. If there is a cost associated with the upgrade

Ethernet's Use in Building High Speed SONET Free Internet Infrastructure

COOK Report: Would you begin to explain then how gigabit and ten gigabit Ethernet and even fast Ethernet is being used as a transport technology to replace high priced SONET equipment in MANs and WANs?

Frazier: Remember that in the case of the enterprise it is not so much the cost of SONET as what you have to pay carriers using SONET equipment for their services.

COOK Report: What about the need for SONET framing running between cities? I am told you can get that from Cisco routers without having to pay for the expensive SONET equipment, but that you do need it. You mentioned something about Ethernet framing at OC192. Can you go Ethernet all the way when the ten gig work is done.?

Frazier: I would say that most of what you said is true. You don't absolutely have to have SONET framing in order to be able to transit arbitrarily long distances. You must have SONET framing if you want to interconnect with the carrier's infrastructure. If you want carrier "x" to provide with a pipe from New York to San Francisco, they are going to provide you with a SONET pipe. Now the other thing they are going to insist that you use is something called LTE which stands for line termination equipment. This will accept a gigabit Ethernet frame from your router or switch and turn into SONET for you. They call this a "managed service."

When you hear people tell you about a managed gigabit Ethernet service what they are really selling you is a gigabit Ethernet pipe. They take the pipe from your customer premises equipment through an LTE where it is framed, multiplexed and stuffed into a SONET synchronous payload envelope and send it through their photonic infrastructure in a SONET frame. This is a so called managed SONET service.

With gigabit Ethernet it is possible to span distances of hundreds of kilometers by building your own infrastructure. Lets take an example of where a customer wants to link a facility in New York City with one in Princeton New Jersey. You live close to Princeton and as you know there are a lot of large corporate campuses and data centers in that area.

Supposing any of the 20 or more fortune 500 companies on the US one corridor from new Brunswick to Princeton wants to link

their Manhattan offices to New Jersey? I should know what they do because they are all my customers. Princeton is about 100 kilometers (60 miles) from Manhattan. These companies can lease dark fiber from central or downtown Manhattan going to New Jersey either through the Lincoln Tunnel or the Holland Tunnel. They go down the New Jersey Turnpike to Cranberry a few miles east of Princeton or down Amtrak where the main New York to Washington railroad line runs about a mile to the east of the Princeton University campus.

So they can make a link between their offices. If they make that link with a managed 10 megabit Ethernet packet over SONET service, the carrier takes ten meg Ethernet, turns it into SONET, and sends it to Princeton where at the customer premises it turns it back into 10 megabit Ethernet all for the princely price of \$8,000 a month! If they want a managed 100 megabit Ethernet service with fifty megabit guaranteed rate, depending on multiplexed service at any point in time, your service level agreement would say that they could start dropping your packets the minute you pushed over 50 megabits per second. That costs \$15,000 per month. If you want lease a dark fiber pair that will cost out at \$16,000 a month for a year-long contract. Now you basically own that dark fiber and can put anything you want to power it including gigabit Ethernet and soon ten-gigabit Ethernet.

What do you need to put gigabit Ethernet on the glass? I'll use Cisco products as an example, but I won't pretend to tell you that other companies cannot provide you with exactly the same thing - our products are just better. So you put one of our Catalyst 6000 enterprise switches in Manhattan and another one in Princeton. We have gigabit Ethernet interface optical modules which plug into the Catalyst 6000 switch. That optical module itself can span the distance from New York to Princeton over single mode fiber. You need no additional equipment.

COOK Report: And no repeaters or amplifiers?

Frazier: Indeed, none below 100 kilometers. But if you wish to exceed 100 kilometer transmission distance, you need to put in amplifiers. You can start by putting an amplifier in at the New York end and one in the Princeton office, this would actually get you from New York to New Haven and from Princeton to Philadelphia or Wilmington Delaware. If you want to go on further to your Washington office whom ever you are leasing the dark fiber from has got to provide you a co-location facility, in this case, probably in Wilmington. You can work the problem out with your supplier. It's just a

matter of daisy chaining one amplifier to the next.

Another way to do it if you want to get from New York to Washington is to put a layer two or layer three switch in your Princeton office. In such a case some traffic is dropped off in Princeton and some goes on to Washington. What you are doing is building your own network. What makes it really nice it is just like your building network, or your campus network. You are using the same elements every step of the way. You can use switches and routers in the very normal manner that you would use to build your own network.

COOK Report: You are running a pure data network. Any voice or video is encapsulated. If you suffer a fiber cut you are trading perhaps almost two orders of magnitude in cost for not having to buy SONET equipment for 50 milliseconds restoration time for a restoration time of about one second according to Cogent Communications. With a data network and restoration of a second no one will really notice right?

Frazier: You can do a lot better than a second. Keep in mind that SONET restores quickly because it does it at layer one. When we talk about figures on the order of a second, we are talking about using routers to restore at level 3 via OSPF or whatever. But there are plenty of other schemes out there for doing restoration. Cisco has a scheme called HSRP that restores much faster than a second.

COOK Report: How long have you had it?

Frazier: We have had it quite a long time. Cisco talks about it but you won't hear about it from anybody else probably. Our customers certainly use it. Now there is also work going on in the IEEE and the IETF on improving restoration times at layer two or layer three — where ever you would like to do it because we recognize that this is one of the things that the SONET world touts. We have customers who are able to get divergent fiber paths out of Manhattan to New Jersey.

On these divergent paths they can also do something we call trunking. There is the standard we call 802.3AD. In this standard it specifies a way to treat multiple parallel links as one logical link. So if you run one link through the Lincoln tunnel and another through the Holland tunnel, you treat them as one logical link. When both are up and running you get two gigabit per second transfer rate between New York and New Jersey. If one of these is lost, it degrades in less than a second to a link of one gigabit. You would drop a few packets.

SONET does give good protection but at huge extra cost and there are plenty of other ways to get equally viable and much cheaper protection on Ethernet links. You can apply optical amplifiers to fiber and can generate either a gigabit or ten gigabit LAN signal. With cascaded optical amplifiers, a well-engineered cable plant and a well-engineered optical amplification system, you can go hundred, a thousand kilometers or more. You can in other words span an arbitrarily long distance. You can also apply things like dropping in layer two switches or layer three routers. Doing this would allow you to build a mesh network and there is no reason that would keep you from being able to build a gigabit or ten gigabit fiber network that covers the whole United States and not have a single bit of SONET in it. This is what Bill St Arnaud is doing with Canarie in Canada.

Economics of Building Enterprise Gigabit Ethernet Networks

Now lets go back to the economics of this. A managed 50 megabit SONET service costs you \$15,000 a month. Gigabit Ethernet fiber costs you \$15,000 a month. When it comes right down to it the cost of the optical components that you have to get from New York to New Jersey is about the 15,000 cost on one months 50 meg managed SONET service. We are not talking about the cost of the Catalyst 6000 enterprise switch which is about \$40,000 with eight gigabit Ethernet interfaces and you have in each office anyway to manage your office LANs. The total cost depends on how many components you stuff inside and how many gigabit Ethernet ports you buy for the switch.

The list price of the optical components that you need to be able to span the 100 kilometers is \$7000 per port. You need one of these at each end of your link. That is less than what you would pay for managed carrier service. For one month's worth of SONET service you own the equipment and don't have to pay for the service any more.

COOK Report: So who owns the dark fiber that your customers are buying?

Frazier: Most of the dark fiber in the world is not owned by the traditional carriers. My understanding is that if you take the fiber owned by Sprint, MCI-WorldCoM, and ATT you will find that it represents less than 20% of the fiber available for lease in the United States today. The majority of the fiber has been pulled in by companies like Qwest, Level 3, Williams, Metromedia. These guys are offering dark fiber at very competitive rates. Now companies like Morgan Stanley, Squibb, Dow Jones and Johnson and

Johnson are building gigabit Ethernet networks. But so are public school districts! The school system in Freehold New Jersey was able to get fiber from their local utility at something like a dollar per month per mile for their own gigabit Ethernet network.

COOK Report: What was the non Ethernet way of doing all this two years ago if you only wanted sub gigabit bandwidth?

Frazier: Fast Ethernet and FDDI were able to go up to sixty kilometers with single mode fiber. Some people did put in 100 megabit per second Ethernet in 1998. But the availability of dark fiber was then not what it is now. Gigabit IP over dark fiber became available about two years ago. It is now getting much more of a "buzz" because of the increased availability of dark fiber and because all the equipment manufacturers have developed full lines of equipment to power these new networks.

If you needed sub gigabit bandwidth, your only choice was to get a managed SONET service from your local carrier at OC3 or DS3 or whatever speed they deemed appropriate. They could also buy a managed Ethernet or fast Ethernet service. These services get converted into SONET by an LTE. The ILECs were offering managed FDDI, managed fast Ethernet, ATM, Packet Over SONET services which were priced outrageously high.

Three or four years ago if you owned some long distance stretches of fiber, you could do packet over SONET (POS or SONET framing) and save some money. The motivation for doing POS was to have the most efficient way transmitting IP via the SONET frames of a carrier.

SONET Needed only for Interfacing with Carrier Equipment

COOK Report: But now if you do POS your only reason for doing so is your need to interface with carrier equipment?

Frazier: Yes. The carrier wants it. I don't. I don't get anything from it. I think the carriers will eventually get rid of it. Right now it is a convenient and useful tool. Its most important aspect is that it is a generic outgrowth of the demands of circuit switched voice traffic. It is easily multiplexed at various OC rates. You don't have to go in one leap from OC3 to OC48 but can buy just the appropriate fraction thereof. It has nice facilities for fault diagnosis and management that they are comfortable with but it is something that the economics of the Internet will force them to do without.

COOK Report: Then, with the completion of the Level 3 and Williams build-outs, to name just a few, the fiber owners are very happy to sell long haul dark fiber to your clients?

Frazier: Absolutely.

COOK Report: But what do I have to do to construct my own fiber network?

Frazier: I am not an expert in this area, but from what I have found it is not very hard to find fiber in any market. Certainly people are able to find connections within the same state. In counties within states you can find dark fiber. I think it is getting down to the point where you can connect any two towns of 100,000 population or more. My impression is that within a couple of years just about any town will be connectable no matter how small. The challenge of course is if you are an enterprise and you want to be able to span great distances, you can't own all the facilities necessary to span those distances.

I believe that one of the next business models to emerge will be alternatives to the carriers. These folk will be very happy to put in equipment to generate a lambda or wavelength on a fiber and they will say: we don't care if you use SONET or not.

COOK Report: Suppose then that I wanted to light a fiber pair between Manhattan and Princeton at a fast Ethernet speed. What are my most basic options?

Frazier: For only 100 megabits you would need to get either FDDI or fast Ethernet optical components to plug into a commercial switch. The components made by the switch manufacturer will typically go to only about sixty kilometers. You can, however, buy third party equipment that will reach the full 100. It takes the form of a transceiver box that sits in your premises in New York and in Princeton. It will boost your optical signal. It acts like a translator, taking a weak optical signal in and putting out a strong 100 kilometer distance spanning signal. The cost for a pair of these components would be in the \$5000 range.

Economics of Building Your Own WAN

Now beyond fast Ethernet, the next logical step is gigabit Ethernet. Many of our customers are taking a slightly different alternative and buying a metro DWDM device. These devices can put multiple communications channels down the fiber pair with which you started off. Moreover they can accept input signals at a variety of protocol framing rates. You can put ten megabit Ethernet into it, 100 megabit Ethernet, ESCON, Fiber channel, Gigabit Ethernet, SONET and you can have up to 32 chan-

nels all simultaneously travelling down the same fiber. You theoretical maximum aggregate traffic using such a device could be as high as 80 gigabits per second.

A system like this would assume a cost of \$12,000 per wavelength or channel on each end of the link. So let's say \$25,000 per channel. Remember that this is a one time cost for up to 32 channels over a single fiber pair. You may populate the box by buying sets of channel cards to install at each end of the link. The whole point is that you can scale this concept of using dark fiber from one wavelength of bare bones Ethernet at five thousand dollars to gigabit Ethernet at 15 thousand and you could go the next step of ten gigabits possibly at a 50 grand cost. You can do this for certain all the way to the ten gigabit Ethernet level by the end of this year. You then can go the next step of DWDM where you can put several lambdas of 10 gigabits on the fiber pair.

Now on the DWDM at 25 thousand dollars per channel, you are left with a maximum of OC48 in speed. I would expect that the OC192 or ten gigabit per second dense wave division multiplexing channels are considerably more expensive than 25 thousand dollars a piece. My understanding is that they can range as high as \$250,000 per OC192 wavelength. To sum up however you have a relatively low price of \$5000 to a higher price of \$50,000 where you can span three orders of magnitude in dark fiber bandwidth increase.

COOK Report: Wouldn't this have to scare the hell out of the carriers? Because doesn't it limit severely how much they can get from the growing data business.

Frazier: If they don't themselves get into the dark fiber leasing business what you say may be quite true. They can certainly help serve the portion of the market that will never build out its own equipment from coast to coast. As I said a moment ago, I think that the advent of dark fiber is going to enable the emergence and growth of a lot of what I call alternative carriers.

Source for Our Earlier Statements on January ITU IETF meeting.

Earlier this year we mentioned an important IETF - ITU meeting in Geneva. This meeting was an International Telecommunications Union sponsored an IP-Telecoms Interworking Workshop focusing on Numbering, Naming, Addressing, and Routing that was held in Geneva from 25-27 January 2000. The web pages detailing the meeting and on going follow up activities are available at <http://www.itu.int/ITU-T/ip-telecoms/ip-telecoms.htm>

Complex issues in IPv6 Addressing Allocation and Administration

Editor's Note: What follows is a discussion of IPv6 address assignment and cost of maintenance issues. The still undecided questions are huge. If ICANN survives the address cost issues will be directly related to its budget.

On August 3 **Dennis Glatting** wrote to the IETF list: I've been thinking about the issue of ARIN fees from last night's plenary and arrived at two philosophical questions.

I run my business out of my home and my DSL link is an important part of my business. About six months ago my ISP started charging me a \$20/mo. fee for my /27 because "ARIN is now charging us." I am unhappy about this fee but I understand its motivation — conservation of IP space, though I believe fees do not really effect the true wasters of this space and the fee, or as it is called in some circles, a tax, is probably misguided. Nonetheless, with IPv6, I naively hoped, until last night, the conservation of space issues would go away, and thus the fees. Big duh!

Brian Carpenter: To be completely clear, the conservation argument has gone away as a constraint of equal importance with route aggregation, and in IPv6 we can give route aggregation much more importance than conservation in the allocation policies. But an allocation service still costs money and somebody has to pay for it. Some of the early deployment scenarios for IPv6 are intended to bypass this issue; for example see [draft-ietf-ngtrans-6to4-06.txt](#).

Glatting: If we look at today's marketing hype and think forward a bit there is a thrust to "Internet enabled" appliances, such as dryers, ovens, and stereos. Assuming ARIN fees persist, my first philosophical question is whether any consumer of these appliances MUST periodically (e.g., monthly) drop coins in the ARIN fountain?

Thinking laterally, the reserved port space (<1024) is tight. Using the same IP space conversation logic, should fees be charged to conserve port space? If so, my second philosophical question is what is our role, as protocol designers and IETF volunteers, in creating, what is slowly becoming, an Internet consumption taxation model? Imagine for a moment the effect of a fee against the allocation or use of port 80 or 443, maybe even port 25 or 53.

Magnus Danielson: Well, who would pay for allocating a new service? It does not matter if it is 10 or 100000 servers offering something on an allocated port. It is allocated. There is also no good way to get ports back, since how can you

Continued on page 21

An Ericsson Perspective on Global Wireless Standards and on Chinese Wireless Market Credits Broad Asian European Market Penetration to Government Pressure for Regional Standards and Interoperability

Editor's Note: Gary **Pinkham** is Vice President for Investor Relations and Business Development at Ericsson. He has been with Ericsson for almost fifteen years. Before that he had worked for a number of other equipment vendors and network operators for a total experience of about twenty-five years in the telecommunications industry. His background is centered in network planning and engineering, project management, strategic planning, sales and marketing, and account development. We interviewed Gary on June 22, 2000

COOK Report: In the commercial wireless world here is much talk now about third and fourth generation systems. For those of us without much knowledge in this area would you please elaborate what is going on.

Pinkham: If we talk about wireless generations, when first generation is mentioned, the reference is to analog mobile telephone service. Second generation was basically digital mobile telephone services, with limited data capabilities. And for the third generation we talk about wide band or broadband multi media services.

COOK Report: Okay, but would you summarize the contenders and their market positions in the second generation?

Pinkham: In terms of companies or technology?

COOK Report: Both. As well as markets and licensing and so on.

Third Generation Services and Migration Paths

Pinkham: If you look at global development, strictly Western Europe in particular, Scandinavian regions and parts of Asia Pacific, like Japan, Hong Kong and Singapore, have achieved a very high penetration level of digital and mobile telephone service, a penetration that is in the range of 60% of the population.

COOK Report: And the reason for this is that the governments in the respective countries settled on a standard?

Pinkham: Yes, they chose a national stan-

dard, and then agreed on inter-operability with other countries.

COOK Report: And the standards generally chosen were?

Pinkham: GSM.

COOK Report: And is GSM an ANSI standard?

Pinkham: There's a version of GSM that is an ANSI standard. PCS, here in the States. And it's also an ITU standard.

COOK Report: There were two other things that were ANSI standards, I believe, yes?

Pinkham: In the U.S., there's actually three systems that are standardized. GSM and CDMA are two. The third is called TDMA, which is also known as IS136.

COOK Report: Those the ANSI ones.

Pinkham: And then you have cdmaOne, which is also known from a standards perspective as IS 95.

COOK Report: And would you help me understand why although someone like Guilder touts CDMA as, quote, a "superior technology," from your point of view, it hasn't been such a technology. At least it doesn't have the market penetration that you would expect. Why is that?

Pinkham: Well, I believe that wireless is about how attractive you can make the applications and how you can create an economy of scale in terms of volume market acceptance. And in the narrow band world, that is second generation, there's no really significant difference between the capabilities of GSM, TDMA or CDMA. It's all more or less the same or equivalent, if you will.

COOK Report: Why?

Pinkham: Because it's all the same bandwidth, and it's all the same sets of Services. Basically, it's just mobile telephone service. Looking at the field you ask which one has the biggest volume, from which you can get an economy of scale.

COOK Report: But you also were linking it to applications? Would you expand on that

a little bit?

Pinkham: Again, the consumer doesn't care what technology is. Consumers just care about how the applications work, how good the quality of the service is and so on. This is because, at the end of the day, no matter which technology you use in the second generation, they're more or less equivalent.

COOK Report: And with second generation we're generally talking about mobile phone service, in other words about cellular phone service.

Pinkham: And a very sort of limited datacom capability. For access to the Internet and what-not. Then when we go to third generation or 3G, there you have a different time frame, so you can apply a different technology. And there's three technologies that have been specified by the ITU for third generation.

COOK Report: What is the different time frame?

Pinkham: Because it's ten years later. The second generation was specified in the early 90's. So now we have a different technology basis that we can leverage. And the ITU has recognized three primary standards for 3G. First is Edge which is a TDMA technology that delivers 384 kilobits of packet data service.

COOK Report: And who's primarily invested in that or offering it?

Pinkham: Ericsson is one of the leaders in that area. And Edge is intended to be the migration path for operators that do not get additional radio spectrum. So it would be the only path for TDMA operators here in the States or GSM operators overseas.

COOK Report: In other words, you're looking at some backward compatibilities?

Pinkham: You just add-on to the existing network and it fits into an existing licensed spectrum and so forth.

COOK Report: In other words, you're taking your technology set and having to build onto it and expand its capabilities.

Pinkham: Yes, but it also gives you an improved data capability on the order of twenty-five times what you have in the sec-

ond generation.

COOK Report: When you are going into a third generation that does not have to be used by people who already have licensed spectrum, what happens? What are the differences there?

Pinkham: Edge primarily gives operators which do not have additional licensed spectrum, an upgrade path to 384 kilobits of multimedia service capability. You have broader bandwidth. 384 kilobits is more than you get off your cable modem at home today. That puts it in some perspective, right?

COOK Report: Yes, it's certainly as good as DSL or a lot of DSL's and the cable modem people advertise more, but whether they deliver much more is another question.

Pinkham: Yes. We are talking the effective rate. So this is what the existing operators do with TDMA and GSM. By late next year we can expect TDMA technology and the GSM technology merge into one common technology called Edge.

COOK Report: Okay, and you said certainly Ericsson is involved in doing this. Who else?

Pinkham: I'm sure Lucent, Nokia, Motorola, Nortel and everybody else is doing it.

COOK Report: Qualcomm?

Pinkham: No, Qualcomm does not have TDMA technology. It is a CDMA player. Now the existing CDMA operators, who are using what's called *CDMAONE*, have an upgrade path to something called CDMA 2000. And there the equivalent to Edge is something called 1XRT. And 1XRT gives you the equivalent capability of what you get with Edge, that is around 384 kilobits of packet data capability. And that's in the same timeframe as Edge, as well. Within the next 18 months to two years, we'll begin to see this faster technology start to be deployed.

COOK Report: Is this like Ford on the one hand and Chevrolet on the other? How do the people who are making the phones compare them or does it make sense to compare them?

Pinkham: They're more or less the same. It's just two different variants, whether it's VHS or Beta is how you might think about it.

COOK Report: Okay, that's a helpful way of describing it. So please continue.

Pinkham: Those operators who do get additional radio spectrum in newly licensed frequencies would probably deploy a technology called wide-band CDMA.

COOK Report: And the new licenses would be in a completely different part of the spectrum?

Pinkham: Exactly.

COOK Report: What spectrum are we talking about?

Pinkham: Two gigahertz.

COOK Report: As opposed to the earlier ones?

Pinkham: They were at 850, 900 megahertz or 1800 or 1900 megahertz. Wide-band CDMA gives you broadband or wide-band capability that delivers from 384 kilobits up to 2 megabits. But it won't stop there because there's other enhancements coming to wide-band CDMA that could lift it up to 5 megabits or more in terms of bandwidth. Wide-band CDMA can also be used as a longer-term upgrade path for the Edge installed base as well as for the CDMA 2000 installed base. That would be for the longer term, though. At the at which we upgrade to the five megabit capability.

COOK Report: And again, the reason for upgrade is to make them more data-capable, more multimedia, more Internet capable?

Pinkham: And also more capacity for the voice services.

COOK Report: I'm a little naïve in some of this. I was thinking in terms of just a single user. When you say more capacity for voice services, how does that affect the single user?

Pinkham: Because the more users you have, the more users you have within a particular area, and every user consumes some of the radio spectrum available to that area. So you need more capacity to handle more users. And the trend certainly is such that the number of mobile users are going to exceed the number of wire line users by a significant amount.

COOK Report: But you're talking now about a CMDA or a TDMA technology?

Pinkham: Both.

COOK Report: But I thought that, because of the code division multiplexing Capabilities, CDMA could theoretically at least do an easier job of getting more users into a given amount of spectrum.

Pinkham: This goes back to these kind of games that technology people can play. To a certain extent, that could be true, but then when you get a fully-loaded network, the TDMA has the same capacity as the CDMA network.

COOK Report: Why is that?

Pinkham: Well, in CDMA, everybody broadcasts at the same time, so the more users you have, the more they interfere with each other and then you have to have smaller and smaller cells to have better coverage for the users.

TDMA you have a take a different approach. There you have what's called pico cells or hierarchical cell structures. You have pico cells, regular cells and macro cells, so it's an umbrella kind of an effect. With CDMA you just segment the cells and make smaller and smaller cells. So when you get through all that process, you still at the end of the day have more or less the same capacity.

COOK Report: So what would be the next part of the evolutionary step?

Pinkham: The reason why wide-band CDMA is different is it uses the whole spectrum and uses a different modulation technique.

COOK Report: Different from the earlier CDMA or different from TDMA?

Pinkham: Both. It uses much more advanced encoding schemes. It uses a CDMA technique, but a different implementation approach, if you will. So that's why you get the significant benefit in terms of the same spectrum with wide-band CDMA. You can get two megabits of service, compared to CDMA 2000 or to TDMA Edge, where you can get 384 kilobits of service.

COOK Report: And is it safe to assume that all the chips and all the technology issues are fully in hand? You did say that you'll be rolling this out within the year?

Pinkham: We start deployment of the wide-band CDMA in Japan. We start shipping at the end of this year and that system goes into service during the second quarter next year.

COOK Report: So it's really done, then, in every respect, would that be accurate to say?

Pinkham: Well, relatively speaking. It's the early phase of a new technology, so certainly we have to continue to develop. As I said then, we'll see wide-band CDMA starting to be deployed beginning in Japan second quarter next year and then you would see other markets, such as Finland, United Kingdom and so on following on after that. And then the 3G, in terms of wide-band CDMA, really ramps up fully during 2002 and early 2003.

When we look at Edge and CDMA 2000, we find that this really begins the end of 2001, and early in 2002 and comes up and

begins to ramp up in the same timeframe. All the same wide-band CDMA will be the first 3G technology out.

Ericsson and Qualcomm

COOK Report: People say Ericsson and Qualcomm have been going at the third generation standard from a different point of view for some time. You said you didn't want to make the description of Ericsson versus Qualcomm, but, if that perception is certainly out there, some people consider it to be something of a battle. Also it seems that you've had some success in China lately that seems to indicate that your version of third generation is winning over CDMA? Is it winning because CDMA's obsolete? If so, why? And what exactly is IMT 2000? Where does that fit in?

Pinkham: IMT 2000 is the ITU definition of 3G. First of all, a comment relative to Qualcomm, we had a dispute over CDMA patent infringements, where we felt they were infringing on our patents. And that was settled last year, so that we see these areas of the technology as not being an issue now. Indeed we're cooperating in areas of further CDMA development.

COOK Report: So how that changes things is that Qualcomm has an installed base of on the order of 50 million, while there's 300 million GSM sets, which use time division multiplexing?

Pinkham: That's right.

COOK Report: And the idea, then, is that, fair enough, you've got two technologies with two market shares and you don't have to consider them, legally at least, at war with each other and they can just go ahead and continue to develop and compete in the market in general.

Pinkham: You must remember that, Qualcomm does not make hardware anymore, we bought the infrastructure business. We're a full systems supplier. We supply all the technologies, so we really don't care which technology an operator wants to use. We focus on what's the best way to support their business development. So we provide *CDMAONE* technology, we provide GSM, we provide TDMA, we provide PDC. We provide all the 3G technologies as well.

COOK Report: Qualcomm is making some technology, is it not?

Pinkham: They make the chip set. And they also develop intellectual property rights around CDMA.

COOK Report: And where does their CDMA chip set fit into a device in terms of

what you make?

Pinkham: We're making a CDMA handset and there we'll use the Qualcomm chip set.

COOK Report: In other words, this is like making the entire component as opposed to just a CPU that's used in a PC?

Pinkham: That's right.

COOK Report: But are some people are claiming that CDMA is obsolete?

Pinkham: First of all, I haven't heard anybody claim CDMA to be obsolete. I mean, CDMA is a good technology. We're using it as the basis for our third generation, which is a form of CDMA called wide-band CDMA. Both CDMA 2000, and *CDMAONE* are valid technologies and is being deployed today by some operators. And, of course, TDMA is also being deployed.

So, again, let's not make it a technology issue. Now, when I was recently in China, it was clear that China is predominantly a GSM market today.

Third Generation Allocation Policies

COOK Report: A couple of quick questions before getting into China. Could you summarize for me some of the licensing and the bidding and the perceptions out there? All I've heard is billions and billions are being bid, enough so that the expense involved may hinder some of the companies from building out infrastructure as rapidly as they would like. What is your general assessment of the licensing and the spectrum bidding process. Is it behaving the same worldwide? Is it different in Asia from Europe and from the U.S.?

Pinkham: It'll be different in different parts. Some places it will be a bidding process, other places it'll be what's called a beauty contest, where the operators will talk about how, by giving them the license, they will be able to do more for society or the country, or economic development or what have you.

COOK Report: But it's for use for a given portion of spectrum, and also use in a specific geographic area, is that right?

Pinkham: That's right and in most markets, other than the U.S., it's for a specific standard. Like in Japan, you can either go with wide-band CDMA or CDMA 2000. And the two largest operators have chosen wide-band and the third operator has chosen to go with CDMA 2000. In Europe, they've all agreed to go wide-band CDMA. In U.S., I guess

it's open for whatever, continuing as it has in the past. You get the spectrum and you can do whatever you want to with it.

COOK Report: How would you evaluate, from the point of view of someone who presumably wants to develop the technology as rapidly as possible and expand the market for the technology as rapidly as possible, the spectrum allocations in Asia and Europe, and in the U.S.. Where is it being done right and where is it not being done so well and why?

Pinkham: I don't know about right and wrong. Perhaps, better or worse?

COOK Report: OK. Better or worse.

Pinkham: I think in Asia and Western Europe, they have a good path, because they have a good spectrum plan. They allocate the spectrum and specify technology to be deployed. That way they ensure interoperability and a mass market.

COOK Report: In other words, in such-and-such a spectrum, the technology that you can use is so-and-so.

Pinkham: Exactly.

COOK Report: And if you're going to use code division multiplexing, you use it in this area of the spectrum.

Pinkham: That's part of the concept. So then they set aside the spectrum to ensure that it's the same spectrum that other countries would use. Consequently, you have this kind of interoperability or transparency. It's the same system everywhere you go.

In the U.S., they haven't quite done it that way and now when you take G3 for where Japan and Western Europe have allocated spectrum for wide-band CDMA, you find that in the US, this spectrum is already being used for PCS. So now what does the U.S. do for third generation? It has to go a different frequency spectrum.

COOK Report: And that means that if I want to talk to somebody in Europe, I've got some compatibility problems.

Pinkham: So then you've got to have some kind of multi-frequency telephone to enable that. That's one of the issues.

COOK Report: And of course it's more expensive to produce a multi-frequency phone and more expensive to buy such a thing.

Pinkham: Of course. Latin America and Brazil have recently taken some steps to clean up their situation. Yesterday, (June 21) when Brazil announced their spectrum allocation for PCS, they chose frequencies, in

such a way that they also could set aside spectrum for the 3G stuff, in such a way that they won't come into the same situation as we have here in the U.S. They followed the European model.

COOK Report: Did we get a little off track in the U.S., what, three, four years ago? Or are we back on track now? And if so, why?

Pinkham: I don't know, you'd have to talk to the FCC about why they allocated the spectrum the way they did it and how they allowed the different technologies to be deployed and what view do they have in terms of the development of the U.S.

COOK Report: So it's a matter of the education of the regulators, in part? Is that a fair statement?

Pinkham: Yes. And also going back to special interest groups trying to promote one national technology over what is perceived as some other national technology instead of looking at the total picture.

COOK Report: And the total picture being global interoperability for travel and communication?

Pinkham: Exactly. You'd better look at the global market and not at the national market. And you better cut down on the influence of the special interest groups.

COOK Report: Is PCS a special interest group? What would be a special interest group?

Pinkham: No, if you look at the different technologies - CDMA, TDMA, GSM, PCS and so on, there's lots of different technologies that were promoted. So instead of the government taking a stand and saying we believe Technology X is the best for both the short-term as well as the longer-term and that's what we're going to standardize on, in the United States our mantra has been let the market decide.

COOK Report: Indeed, the way our attitudes have been, that's the government picking a technology winner and a strong attitude has been, oh, you can't do that. And you're saying, okay, if you can't do that, and we didn't do that, but here's the situation it's left us with.

Pinkham: Exactly. We're so far down the first path, that I don't know how we recover from that.

Future Dominance of Wide Band CDMA

COOK Report: That's very interesting, be-

cause if you try to take a global view of this and I suppose the North American market is still the largest, but if you put Asia and Europe together, those are surely larger than our market.

Pinkham: No, Europe is bigger than North America.

COOK Report: Already? Alone?

Pinkham: And China alone will soon be bigger than the U.S.

COOK Report: And is it correct then to assume that not only have you European standards and you're going to have a Chinese standard now and the Japanese, but to what extent are the Asian standards compatible with all the other Asian standards and compatible with the European ones?

Pinkham: The wide-band CDMA — that is the standard.

COOK Report: In all those areas?

Pinkham: Definitely. And a compatible frequency as well.

COOK Report: So where does that leave everything?

Pinkham: Just look at the demographics, you can see why wide-band CDMA will be the predominant technology for third generation. Followed by Edge, because Edge is the upgrade path for the installed base of GSM.

COOK Report: But, again, when you say look at the demographics, lead me by the hand and the demographics for wide-band CDMA are what?

Pinkham: Okay, all of Europe is standardized in wide-band CDMA. Japan is standardized in wide-band CDMA. China will go for wide-band CDMA. Count the people and you'll see how that goes.

COOK Report: But you've got GSM still out there?

Pinkham: GSM because of the upgrade path to Edge, will also be predominant, yielding Edge as a predominant third generation technology.

COOK Report: So the GSM users in third generation will be Edge users?

Pinkham: That's right.

COOK Report: And what users will be third generation broadband CDMA users?

Pinkham: All the new licenses.

COOK Report: And all the current, earlier versions of CDMA?

Pinkham: They will probably go to CDMA 2000.

COOK Report: CDMA 2000 and broadband CDMA, it sounds like you're talking about them separately.

Pinkham: They're two different standards.

COOK Report: And the one that from the point of view of demographics and market that is positioned much better is broadband CDMA, is that right?

Pinkham: That's right.

COOK Report: And Ericsson in broadband CDMA is making everything you think should be made?

Pinkham: Yes, but we also, we're making a complete line for CDMA 2000, as well as for Edge.

COOK Report: And Qualcomm at this point is more in chip sets and you said some intellectual property, I gather you mean software.

Pinkham: Yes, the patents.

COOK Report: In other words, you're making complete CDMA equipment and cell phones and you have some Qualcomm chip sets in there, but are you in effect licensing some Qualcomm patents when you make the complete devices?

Pinkham: But we have patents as well. To be clear, Qualcomm does not have a monopoly on CDMA patents. We have a strong patent position as well.

COOK Report: So for any given product, presumably, you have a choice whether you use vanilla or chocolate.

Pinkham: No, it's always a mixture. It's a cross-licensing agreement, so depending on which product, we maybe use more or less Qualcomm patents and more or less of our own patents. That would determine royalty stream. Same thing for Qualcomm, they'll use our patents as well.

COOK Report: And using a mix of these patents gives you a product that performs better in a given market or is more economically viable?

Pinkham: No, the mixture is determined by whatever's required for the standard. The standard is a mix of intellectual property rights are represented as patents, and number of chips.

COOK Report: I see. Are there any remain-

ing areas, in your opinion, of spectrum licensing or standards or any other unresolved issues from the point of view of either technologies or geography?

Pinkham: No, I think it's sort of settled out and it's just now to work it through the process to get everything completed in time. So all the outstanding issues are pretty much resolved by now. We will have version ABC or Release 123, however you want to look at it. So the earlier versions will have a slightly different capability than the later versions, but then they'll be upgraded continuously. I think we have enough planned in terms of starting the deployment and then having an upgrade path.

Wireless in China

COOK Report: What happened in China? If I recall, the Chinese had signed some kind of an agreement with Qualcomm, perhaps, six months ago? Or was it longer than that? And in the last couple of months, they said they want to reconsider it and they've backed off completely?

Pinkham: You'll have to talk with Qualcomm about what they signed or didn't sign with China.

COOK Report: Well, from your point of view, lead me through the China area. Would this be the last great market to be explored?

Pinkham: I don't know. In China there's basically two main operators in global communications. One is called China Mobile Communications. And the other is called China Unicom.

China Mobile Communications is the largest. Out of a total of 50 million subscribers in China, they have more than 40 million. So China Mobile has standardized on GSM technology. China Unicom, which is a smaller player, they have six or seven million subscribers, they were also using GSM technology, but they were discussing using CDMA technology from Qualcomm just to have something different from the main operator.

Over the time with the WTO discussions and all that kind of stuff, they've continued to deploy GSM, and now at this point in time, it doesn't make sense for them to start to deploy *CDMAONE* technology. They would prefer to stay with GSM technology because third generation, as we talked about earlier, is just a couple of years away. Consequently, they would prefer to stay with the GSM for second generation and then go to a CDMA technology for third generation, rather than restarting today with the *CDMAONE* technology.

COOK Report: It sounds like you're saying this is all perfectly logical and rational.

Pinkham: Yes.

COOK Report: No surprises there.

Pinkham: Not for us.

COOK Report: But something happened in China. I'm unaware of what it is and what the ramifications are. Therefore please, in the short amount of time remaining, give me the short course in what's happened in China. Why is it important and what does it mean?

Pinkham: Well, China is a tremendous sized market. This year they're growing at the rate of more than 2 million new subscribers per month. At that rate, China Mobile Communications will soon be the world's largest network operator.

COOK Report: If you look at what their expected growth should be over a period of time, is there reason to believe that the 2 million a month is going to be pretty sustainable for a while?

Pinkham: Yes, they anticipate, I don't remember the exact numbers, having about 250 million subscribers in five years.

COOK Report: You have pointed out because of some technical regulatory choices that we made, what, four, five years ago in this country that we've gotten ourselves a little boxed out in some of these markets. Is what is going on in China now likely to change this or not?

Pinkham: No, because like I said, China Mobile, which may have 250 million subscribers in five years' time, will be GSM-based and then they will migrate to the next generation with wide-band CDMA technology. At the same time China Unicom, which is a much smaller player has decided to stay with GSM for the short-term. In longer term I think they will go with the CDMA technology. Whether it's wide-band CDMA or CDMA 2000, I don't know.

COOK Report: You say short-term they've decided to stay GSM based while you think they will change longer term. Would you elaborate a little bit as to why?

Pinkham: They'll want to go to 3G. In China they have said they would go to third generation sometime between 2003 and 2005. That's when they think they'll need to upgrade to 3G capability. At that point in time, China Unicom would then decide which way do they want to go, which technology do they want to use. In the meantime, it seems like they'll stay with GSM.

COOK Report: Because, according to at

least one press report, I believe they recently cited CDMA as being in their opinion obsolete. Where are they coming from? How would you interpret that?

Pinkham: You should probably talk with somebody that's more familiar with the politics in China, because this is to a very large extent driven by politics between governments, with the WTO and all of that kind of stuff. Consequently there are some illogical approaches. At the same time, China Unicom is going to have a public stock offering. There's also that game being played.

COOK Report: And apparently there's mention of going with the future CDMA standard. Is this just positioning? Politics? Marketing? Or is it something specific?

Pinkham: Probably all of that. I can't comment about what China Unicom, you should talk to China Unicom.

COOK Report: You yourself have just gotten back from China?

Pinkham: That's right.

COOK Report: How would you describe and interpret the experience you had? Where do you see the things going both within China and the impact of what you saw going on in China, the impact of that on the third generation development in the United States and in the world in general?

Pinkham: Two things: First of all, I think that I was surprised with how well developed the Chinese mobile network is compared to here in the U.S. For example, we took a group up to the Great Wall at nighttime and we all had these GSM world phones. We passed them around and all of the visitors were able to call home while sitting at midnight on the Great Wall of China.

COOK Report: That must have been fun.

Pinkham: And you couldn't go to the Grand Canyon in the U.S. and call China, probably—just to put it in perspective.

COOK Report: Why is that? Why can't you go to the Grand Canyon at this point? Because, again, our patchwork of standards?

Pinkham: In the U.S. it's much more than a patchwork of standards, we don't have it built out nationwide. We focus on competing so the major cities are covered well, but on a universal basis, it's not covered. But the China Mobile Communications estimated that they have, on a geographic basis, 80% coverage. Eighty percent. And a country as big as the U.S.

COOK Report: Amazing. And they've achieved this in what period of time?

Pinkham: About the same period of time that we've been working on it.

COOK Report: Which is about, the last five years or the last ten years?

Pinkham: The last eighteen years.

COOK Report: And, in effect, does this mean that all of a sudden, there is a much more advanced telecommunications market in China than perhaps we ever thought possible?

Pinkham: It certainly could develop that way. Based on international standards, with nothing proprietary.

COOK Report: But, in other words, doing a rollout that has the advantage of being based on a national standard. Presumably, there's a policy, or government subsidies that take it into the rural areas as well? Or do you know exactly what it is that's enabled them, other than the national standard, to get such good geographic coverage?

Pinkham: Up until recently, the operators had been owned by the Chinese, because they didn't have a capital market. So as a social obligation, they're obliged to build out for universal coverage. To build out and thereby, cover the need of the most people. And then instead of fragmenting amongst several competing technologies, they standardized on one and marshaled all their resources for a common goal. The second thing to think of, too, is that in China, they probably are bypassing the wired world and going directly to a mobile world, because you can deploy mobile faster.

COOK Report: The great leap forward and all that, there's been some examples of that in Chinese politics. But the implication of that is obvious. This has been very, very helpful.

IPv6 Issues Continued from page 15

assure that there is no longer traffic on some port.

Neither the servers or the clients can be in a sufficient way charged for the usage of a well-known port in order to achieve the same pressure as you can do with (real) IP numbers. ISPs can naturally charge you for open up traffic to and/or from a certain port, but that does not give the knowledge that you would require. If someone is running some ancient protocol in his/her network using some well known port and would not be communicating it to the outside world, it would still make this port occupied without any money changing hands.

Sadly enough I don't think money is the way to solve conservation of port space, we have to rely on good engineering decisions, and boy, does we

know that these are not reliable :)

Port numbers are as they are and to some degree we have run into a couple of scaling issues with them. As always, set a limit and we will (eventually) outrun it.

Jason Rakers: When household appliances begin becoming IP addressable, I think we will see a move towards assigning an Internet IP address per household (much like today's street address). The household will perform NAT for all devices within (one street address can house many people, not just one).

Anthony Atkielski: I think we'll see IP addressable toasters and washing machines just after we all switch from automobiles to hovercars and from telephones to Picturephones.

Lloyd Wood: But before the Yanks can actually roam between all their own cellular telephone networks?

In any case, there's already: <http://www.margherita2000.com/> 'The first washing machine on the internet. She listens to you and she talks back.' (note the E-MAIL on the machine's controller display.) http://www.merloni.com/eng/ari_digi.htm (using the so-called 'WRAP' protocol) <http://www.allnetdevices.com/news/9912/991203washing.htm> <http://www.telegraph.co.uk/et?ac0111464113065&pg=/et/00/5/4/ecfkit04.html> <http://catless.ncl.ac.uk/Risks/20.88.html#subj9>

(The only recent toaster development I'm aware of is the clear toaster, based on treating glass with some rare earth metal, e.g. http://www.appliancemagazine.com/mm/wt/html/confortec_2000_p10.html <http://www.archimagazine.com/nphil1.htm>)

Mikhail Evstiounin: Wasn't avoiding NAT one of the goal of IPv6? I recall a pretty big discussion here some time ago about NAT and IPv6.

Carpenter: No, what will happen is one IPv6 prefix per household/car/whatever, and therefore no reason for NAT.

Andre John Mas: The factor in all this is cheap firewalls/routers. With these coming down in price they will be installed standard when asking for a cable internet connection, or such. Also these boxes, probably being the size of a phoneset would probably also include a DHCP server for addressing the various devices in the home, much like what is done for Sun's SunRay line of devices. If this does happen I just hope that they offer the option for expert users to configure the devices themselves. The price factor that will make these an option is around USD 100.

Jim Stephenson: Such a product is available already, check out : www.maxgate.net/product_3200.htm

I believe the device does use NAT (:<) but there are other devices in the family, that do IPSEC. I am still waiting for mine to be delivered, sorry if I am a little woolly about it, I haven't had chance to play with it yet, but it looks promising. At the moment the cost is about \$300 USD, but as Andre points out the price will drop eventually. Also I had believed that IPv6 had enough addresses for every person on planet Earth and then some.

But here is an interesting question. Would the address be assigned to a person or a property, on the basis that if the person moved would the address go with them or stay with the house? What happened if the house was destroyed or the person died, would the address be returned to the registry, from where it was issued, to be re-allocated to a new property or person?

Will we eventually get to the point where we are all as unique as our IPv6 addresses, and new born children are assigned an address automatically, when they are born? Maybe George Orwell and Aldous Huxley got it right

Jon Crowcroft: In an ideal world, this would be worked out thru a non linear dynamic pricing model the same way the airlines do differentiated seat pricing We have ipv4 and ipv6 addresses; they have different cost recovery models and different utility functions and different marketing dweeb selling them, and different incentive requirements. So by then, the price would be very very low as the cost of dealing with the possible future admin of assigning more addresses will be vanishingly small....

Tim Salo: Geeks like us care about end-to-end transparency. Refrigerator's don't. Most people are going to buy products based on the functions they perform (utility), not on their architecture. If someone develops a nice household/Internet gateway that does something useful (and doesn't require a UNIX administrator), people will buy it, regardless of whether it performs some politically/religiously incorrect protocol transformations.

Keith Moore: true enough...but they will insist that the product work as advertised. and if households have NATs in them, the number of things that those products can do will be considerably reduced. hence the number of products available, and the utility of a household network connection, will also be reduced.

Robert Elz: Refrigerators don't care about connectivity at all. They don't care about anything at all. Refrigerator owners care about connectivity, and even geeks can be refrigerator owners. Personally I want my refrigerator to have end to end transparency (I'd care less about the toaster, doing anything useful to that when not physically close is harder to imagine).

Salo: And, if IPvX addresses cost money, a lot of households will pay money for devices that

enable them to operate with only one IP address.

Elz: True. However there's no reason to assume that it will be possible to get less than a (very large really) block of addresses, so if you pay for the one, you're going to get a lot to go along with it.

There's no question but that address assignment and management costs, and needs to be paid for, and as with anything, it is the end consumers who pay in the end (everyone else just passes along the costs, whether they're separately listed, or just bundled in an overall price for service). However, with IPv6 there's no longer the need for the rationing push on prices that exists in IPv4. It costs a certain amount to perform an assignment (and record it, etc) - it doesn't really matter how big the assignment is, with IPv6 there should be no differentiated costs for address assignments based upon size.

Moore: True enough, at least if the addresses cost more than the NAT box. but if you have to have a NAT box then the money you spend is for the purpose of making your network less functional. personally, I'd rather get something useful for my money.

Folks who think this is a religious argument aren't paying attention. But the bottom line is that we need to make sure that (a) IPv6 address blocks of reasonable size have near-zero cost (b) NATs aren't part of IPv6

Vint Cerf: Thanks to Dennis Glatting for drawing attention to this question. One of the reasons for fees, of course, is that the Address Registries also have responsibility to support ICANN so they have some new costs in addition to their operating costs (or if you like, their operating costs include support for ICANN).

It is a very good question whether one's internet-enabled household appliances will induce a monthly charges - do you suppose there would be a way to have a one-time charge to "pay" for some number of such addresses - perhaps built into the cost of the appliance (and paid by the manufacturer who "burns" an address into the device - at least the low order 64 bits or something to make it end-to-end unique)?

Please don't flame me for thinking out loud - Dennis' point is a good one and we ought to discuss - perhaps in a smaller group than the whole of ietf announce list!

Robert Book: Vinton's idea has much merit. A scheme to allocate blocks of addresses to manufacturers would be much easier to support than an organization attempting to process individual email requests, or CGI scripted forms from a webpage, or a world-wide DHCP server for Amana (and one for Maytag, etc.) to register a refrigerator. And easier administration should translate into lower cost.

John Stracke: Mmm...that's already there, isn't

it? The low-order 64 bits he's talking about are things like Ethernet addresses, for use in IPv6 autoconfig. (We can't allocate blocks of actual IP addresses to manufacturers, of course, or the routers will go mad. I don't want to wake up in the middle of the night to find a zombie router, with its front panel glowing insanely, standing over me with a Cat5 garrotte and screaming, "Revenge!". :-)

Valdis Kleitnieks: You don't want to assign network addresses on a per-manufacturer basis. Network addresses have to be aggregable. That's why we have ARP because MAC addresses on an Ethernet are assigned per manufacturer. But that's OK, because IPv6 has ARP ;)

Vint Cerf: I had been assuming that the aggregation would take place at the "prefix" level of IPv6 addresses and that the unique component would be in some lower order part of the 128 bit space. If I have a bad model of that, I'd appreciate offline explanation.

Keith Moore: The various proposals to burn IP addresses into devices are naive.

IP addresses identify points in the network topology. Hence they need to be assigned according to topology. If you try to assign IP addresses in some other fashion, you immediately need something else to replace the IP address which is assigned according to network topology. And for various reasons that new thing will have a very similar structure to an IP address.

(For reasons of efficiency, IP addresses are also sometimes used to name, or as parts of names for, lots of other things - including interfaces, hosts, and connection endpoints. but it's the use of an IP address to name a point in the network topology that drives their assignment.)

Yes, you can add a layer of indirection to the network so that a device has a constant 'virtual' IP address which gets forwarded to the 'real' IP address which indicates the current location of the device. mobile-ip does this, for example, and it's quite useful to be able to have a mobile device with a stable IP address. but then you still need a 'current location' address. you also need a home agent which sits on the net at the location corresponding to your stable IP address, and either forwards your traffic to the 'current location' address or issues a redirect to that address.

If my toaster is going to be on the net, the last thing I want is for all of the traffic for that toaster to go to the toaster manufacturer's home agent just so it can get forwarded to my house. I don't want the toaster manufacturer to be able to see how often I'm checking on the status of my toast. nor do I want to have to pay the toaster manufacturer to get it to continue to route packets to my toaster.

Burning IP addresses into devices is a good way to give vendors the ability to control those devices, monitor their usage, and to lock their cus-

tomers in to particular services. not my idea of a desirable state. Personally, I'd rather have the address registries.

Dennis Glatting: It might also be a good way for script kiddies to efficiently scan the Internet looking for a particular manufacturer's device to exploit a discovered security flaw, such as turning off a stove's gas pilot and turning on all burners. If that doesn't sound realistic, how about a cracker inside a manufacturer's systems doing the same; or how about a terrorist?

Andre-John Mas: Though if the devices already use ethernet, then each device would already have its own MAC address, and the IP address would be DHCP assignable. As stated earlier once firewall/routers/DHCP server combos come down in price, preferably below \$200, you will probably see these devices as becoming standard. These devices would provide a level of intrusion security and would be able to assign address to the devices in the home, as is already done in most LANs.

Glatting: Firewalls do indeed provide a level of security but they are also vulnerable to attack and code and configuration entropy — there are many examples of this. Also, if you have a trusted party managing your HVAC system and that trusted part is cracked, a firewall will probably provide no defense.

Andre-John Mas: I suppose what's needed here is some sort of EPROM, that stores an image of the system, and re-copies that image to main memory every day. The EPROM could only be written to by pressing a hardware switch, otherwise it is read-only. Since there are Linux based firewall solutions that fit on a floppy, this sort of approach shouldn't be too difficult

Vint Cerf: For the record, I was assuming that only a lower order unique end-id was burned in and that the high order bits would be dynamically assigned based on connectivity to the global Internet.

Bob Hinden: From the ARIN web site I found that the minimum allocation is a /20 and the fee for that is \$2,500/year (<http://www.arin.net/regserv/feeschedule.html>). Larger allocations have higher charges but the price per address goes down quickly. I would think that most ISP have larger allocations than a /20. A /20 provides 128 /27 address blocks. That works out to \$19.53/year per /27. You said that your ISP is charging you \$20. per month. Not a bad markup! Interesting what happens when rationing starts.

The IPv6 address blocks allocated by ARIN are much much larger, so the price per address for an ISP is considerably lower.

Rick Wesson: the ASO members don't support ICANN on a per block basis, in fact ICANN's Task Force on Funding (TFF) observed that the IP Address Registries operate on a non-profit business model from member fees and should

Executive Summary

GiG E in Montreal, pp. 1-9

We interviewed Francios Menard, Lead Engineer Internet Division of Cable VDN in Montreal. Cable VDN is a Montreal competitive cable TV provider. Menard is designing and implementing a gigabit per second Ethernet data network over the same fiber infrastructure by means of which it delivers cable TV. As far as access to fiber optic cable and a city wide infrastructure of conduits goes, Montreal and Stockholm are possibly the most fiber rich cities in North America and Europe. In Montreal VDN owns two cables, each extending for more than 200 kilometers beneath the city. One cable has 288 strands and the other 144 giving it well over 10,000 kilometers of fiber through out the city.

What we see with VDN in Montreal is the beginning of new business models for start up data network service providers. For customers who do not need to interface their data network with a carrier SONET based network, it is possible to build a gigabit per second data net using gigabit Ethernet as the framing and transport layer. The cost of doing this is extraordinarily cheap compared to carrier SONET based managed data services. This leads to five year contracts for point to point, kilometer long, gigabit per second data services at a thousand Canadian dollars per kilometer per month or less than \$700 per month U.S. dollars. The cost then for a thousand megabit per second Ethernet data network service looks to be about 20% more than the cost of fifty megabit managed SONET data services from a carrier. Twenty times the service for only one fifth more money. This price differential makes it possible to "over provision" most customers who when their data use peaks will have a much higher ceiling of bandwidth on their data lines. If their traffic bursts never hit their connections upper boundary, congestion is unlikely to develop making those pesky quality of service issues on their networks irrelevant.

Of course bandwidth demand will likely grow to meet whatever amount is available. Nevertheless Menard suggests that Gigabit and soon ten gigabit Ethernet data service makes it possible to cease emulating the PSTN with protocols like MPLS serving to route restricted data services over permanent virtual circuits where network applications are designed with the expectation that they will be bandwidth limited. According to Menard gigabit Ethernet makes it possible to design the network with from the point of view of unlimited bandwidth and applications that "adapt" to the bandwidth available to them. For example audio codecs can be used that a

very stingy way if the network is crowded. However the user can shift to mpeg2 and then to mpeg3 when bandwidth is plentiful. He suggests that over provisioning with gigabit Ethernet will give application designers room to experiment in creative ways with new applications which will become attractive precisely because they are adaptive. The large initial amount of bandwidth should be attractive to network MIS managers who won't have to be measuring their network consumption to order new lines on a monthly basis.

As adaptive applications are be designed they must incorporate means of locking out users who refuse to throttle back to lower levels of bandwidth consumption when signs of network congestion appear. The network manager will have to set policy for over all application use since some applications can in the hands of some people suck up most available bandwidth. In a well designed network where application tools have been correctly provisioned such tools should be able to talk to the network management system and smooth out usage patterns. What is driving this remarkable growth in bandwidth is the ability for the first time to go SONET free in a high speed network and the dramatic increase in the speed of Ethernet which is the most thoroughly commoditized and wide spread technology in data networking and therefore the cheapest means of data transport.

10 gig E Standards and Economics, pp. 10-15

We interview Cisco Distinguished Engineer Howard Frazier. Frazier is one of the inventors of ten gigabit Ethernet and a key technical contributor to the standards activity. The draft ten gigabit Ethernet standard is due for delivery next month (September 2000). They have selected two specifications for single mode fiber. One spec (the LAN) is good for cabling up to ten kilometers on single mode fiber. The LAN spec is intended primarily for campus applications and for connection into DWDM transponders. The same spec will be good for hooking a switch or a router to a DWDM box. Ethernet speed here is defined as 10 gigabits per second.

They also selected a specification for going long distances on single mode fiber. These are distances of greater than or equal to 50 kilometers. In the WAN specification data is slowed slightly. It runs at the SONET OC-192 rate which is 9.29419 gigabits per second. They only specification they have not decided on is the one for multi-mode fiber. They have five or six competing specifications that they must narrow down to a smaller number. Cisco has publicly announced pre standard 10 gigabit Ethernet products including dual port cards for its Catalyst 6500 switch.

The Frazier interview validates what we learned from the study of the VDN situation in Montreal. With an abundance of fiber and Gig and soon 10 Gig Ethernet from a data network point of view, the different between MAN and WAN (Wide Area networks) has disappeared. Frazier gives examples of gigabit Ethernet service on leased dark fiber from New York City to Princeton New Jersey, a distance of about 100 kilometers or sixty miles. Fortune 500 companies with offices in New York have bought managed SONET data network services from the carriers in droves. Fifty megabit guaranteed service costs \$16,000 a month from the carrier.

If the same company has leased dark fiber from Manhattan to Princeton, it can get gigabit service (twenty times as much of what it is giving up) to Princeton with no carrier involved. It can do this for the \$15,000 cost of a pair of one gig Ethernet port cards for its enterprise switch. The only other cost is staff and fiber leasing experience. After the first month 20 times the telco services become free because the company uses precisely this equipment to provide its own service campus to campus. It seems likely that the Fortune 500 will be dropping their carrier SONET based services in droves and doing all this gigabit networking for themselves as their managed carrier data network service contracts mature. Frazier has observed several cases. With gigabit Ethernet enterprise networks have a chance to become SONET free. SONET free means much less cost and more effective use of fiber.

The implications of what is in effect another revolution are profound. The only reason enterprise networks would need to use SONET is if they needed to connect to a carrier network that did. Virtually all do. Now cross country long hauls used to require SONET. They don't anymore. Frazier shows how gig and ten gig Ethernet may now be used cross country. In effect data networking doesn't need the carriers any longer. They will soon sit with hundreds of billions of dollars of SONET switches to amortize and the only thing they will find these switches good for is supporting their long distance voice networks that are no longer profitable. We find that the carriers are going to realize that they need to get their SONET off their books as quickly as possible. One way to do that would be to sell off their voice long distance services.

Ericsson's Wireless View, pp. 16 - 21

We interview Gary Pinkham, Vice President of Business Development at Ericsson. Pinkham explains third generation wireless services and their migration paths on a global basis. Services covered are GSM, CDMA,

Continued on page 24

cdmaOne, CDMA 2000, TDMA, Edge, and wide-band CDMA which he thinks will become globally dominant. We asked Pinkham why wireless penetration was greater than in the US. He replied: "in Asia and Western Europe, they have a good path, because they have a good spectrum plan. They allocate the spectrum and specify technology to be deployed. That way they ensure interoperability and a mass market." The U.S. free and open market approach he finds to have left us with patchwork and fragmented services. He compares CDMA and TDMA and Ericsson's relationship with Qualcomm. He finds China which he had just visited to be predominantly a GSM market. China Mobile which is signing up two million new users per month he estimates may grow to 250 million users within five years. China may never install ubiquitous wireline services but leap frog instead into universal wireless.

IPv6 Allocation and Administration, p.15, 22-23

From the IETF mail list a discussion of the need for charging fees for the allocation of IPv6 blocs. We find that even though there are several orders of magnitude more IPv6 than IPv4 numbers hierarchical routing will still be necessary.

Continued from page 22

foot 10% of ICANN's budget. (see <http://www.icann.org/tff/final-report-draft-30oct99.htm>)

If ICANN's budget grows the ASO's responsibility grows proportionally. If the IP Address Registries (ASO Members) are doing cost-recovery then its expected the price for membership will increase or the cost for a delegation will increase. It is reasonable to assume that an address maintenance fee will eventually passed along to the consumer just to maintain the 10% ICANN maintenance.

I don't want to see IP Address space charges, no matter its version, induce a monthly charges on the end-user side. Just this year the ASO's 10% responsibility will amount to \$428,000 USD. If you don't think that ICANN's budget will grow understand that ARIN's 2000-01 budget is 5.2M and ICANN's budget for 2000-01 is 5.0M. If you want to keep ICANN out of your pocket, ensure they stay lean and focused on technical administration of Assigned Names and Numbers and not inflating its own self worth.

I would prefer having ICANN set the ASO funding requirements in such a way that it did not encourage ASO members to pass on the ICANN maintenance fees in a per assignment basis. In short Managing ICANN's budget can have the largest impact on the costs of future IP Address Registry operations. I'd be happy to carry this conversation on in another forum, just let me know where that is.

**Battle for Cyberspace-
How Technology and
Political issues May
Affect your Internet
Venture cost of \$695
and now available.**

Subscription Rates

Choice of either ascii or Adobe Acrobat (PDF) format 1. Individual; College or University Department; or Library; or Small Corporation - \$250 2. Corporate - (revenues \$10 to 200 million a year) - \$350 3. Large Corporate- Revenues of \$200 million to \$2 billion per year - \$450 4. Very Large Corporate- Revenues of more than \$2 billion per year - \$550
Site License: The right to distribute ascii and PDF via email to all employees of corporation. 5. Small corporate: \$450 6. Corporate: \$650 7. Large Corporate: \$900 8. Very Large Corporate: \$1150 .
Site License Distribution via intranet web site \$400 a year additional. See www.cookreport.com for more detail

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

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