TECHNOLOGY TRENDS

Will Bandwidth Ever Be Too Cheap to Meter?

Michael Weingarten and Bart Stuck

Probably not, and the alternative—usage fees—is equally undesirable. But there may be another way.

Since the publication of George Gilder's *Life After Television* and Nicholas Negroponte's *Being Digital,* it has become an article of faith among Netizens that by using high-capacity fiber and wireless connections, Internet bandwidth someday will become too cheap to meter. Against this, however, there are dissenting views.

The Economist View: Tragedy of the Commons

The first comes from a group of economists and technologists who participated in a 1995 MIT symposium on Internet economics, and published their results in a book called *Internet Economics* (MIT Press, May 1997).

A recurring theme in the book is the belief that the Internet is doomed to over-congestion unless we meter and bill for usage. These economists do not believe we will realize the Gilder/Negroponte dream (see Figure 1)—especially considering the current Internet "all you can eat" pricing model. For the economists, giving people unlimited access to a finite resource sets up a classic problem called the "Tragedy of the Commons," in which people overuse a free good and render it useless for everyone.

The economists' solution is to charge for marginal usage. They called for real-time metering implemented through Quality of Service (QOS) header information included in the next-generation IPv6. However, since 1995, the potential for widespread IPv6 adoption has become more remote; unless these prospects improve, some other QOS mechanism would be required to provide real-time metering. While the economists' critique of Gilder and Negroponte isn't affected by these QOS implementation problems, their alternative proposal is, as we'll see in greater detail.

The Qwest IPO

A second rebuttal to Gilder and Negroponte can be seen in the recent Qwest IPO (June 25, 1997). Qwest is installing a high-capacity intercity fiber network that in principle could make the Gilder/Negroponte vision "work."

FIGURE 1 Economists' View of Congestion

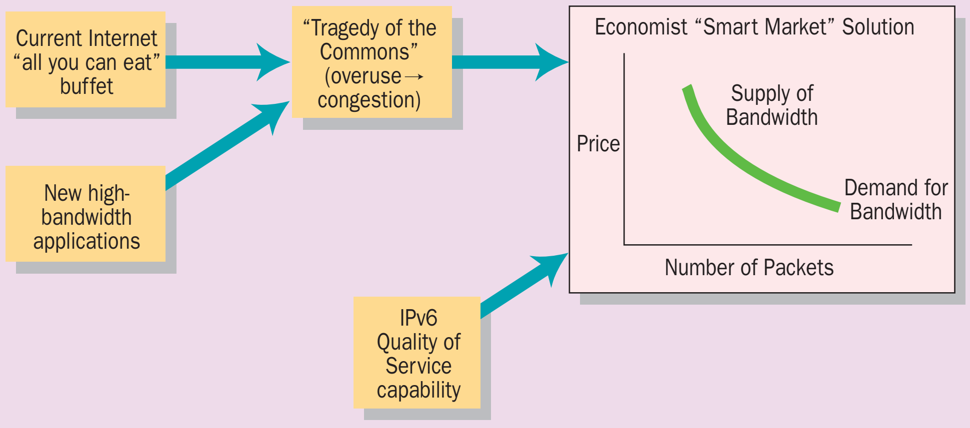

Michael Weingarten (michael_weingarten@monitor.com) is a telecommunications consultant with Monitor Company (Cambridge, MA).
Bart Stuck (bartstuck@aol.com) is president of Business Strategies LLC (Westport, CT), a network computing and telecommunications consultancy.

We don’t foresee the kind of abundant bandwidth that Gilder and Negroponte anticipate

However, Qwest management is not planning to provide “too cheap to meter” capability. Instead, the company plans to deliver bulk bandwidth to carriers in order to alleviate a current shortage in high-capacity long-haul pipes. If end customers were not willing to pay for such transport capacity, Qwest’s buildout would make no sense from a business perspective.

This is echoed by comments from Jack Grubman, Salomon Brothers telecom analyst (and Qwest’s lead IPO underwriter) in December 1996. “We need more bandwidth,” he said. “There are capacity shortages everywhere.… The U.S. long distance backbone network…is a scarce resource.”

The Outstanding Issue

The above arguments point to a dilemma regarding supply versus demand. Do we believe supply will be scarce? Or is there a Gilder/Negroponte-type technological solution that avoids shortages?

We don’t foresee the kind of abundant bandwidth that Gilder and Negroponte anticipate, but we believe that a Balanced Pipe/Expected Minimum Strategy approach, described below, can avoid the economists’ solution of metering every packet across the Internet. We will need to allocate, but if we are clever enough, we won’t need to meter.

Starting Point: Prospects for a Supply-Side Fix

The Gilder/Negroponte supply-side fix involves developing sufficient capacity to cover even major demand increases. It is based on two assumptions:

- Fiber-based and wireless transport capacity is effectively infinite.
- Continued Moore’s Law improvements in transport and switching will support any prospective demand scenario.

We don’t see this in the cards. True, we are likely to see substantial supply-side improvement resulting from low-cost parallel processing. Commercially available SONET products now run at OC-192 (9,920 Mbps), which is 64 times the OC-3 rate that became available only five years ago. There are similar opportunities for improvements in electronic and optical switches, also using parallel processing.

On the other hand, there are three reasons to doubt the prospects for a complete supply side fix. First, nothing like infinite capability is deliverable given likely developments in these key areas:

- **Transport-Fiber—**If we wanted to give every person in the U.S. a T1 line for continuous simultaneous usage (not that much per person, by broadband standards), we would need a network with aggregate capacity of 409 terabits. Even at OC-192, this would require 41,000 fiber strands—or 427 Qwest networks.
- **Transport-Wireless**—Alternately, if we wanted to rely on Gilder’s “ethersphere” concept (ignoring FCC licensing limitations), it would require 100 teraHertz of wireless bandwidth at a coding rate of 4 bits/hertz—far more than the entire available electromagnetic spectrum. Even with cellular reuse, to provide T1 to each nonrural person living in an average 3,250/square mile population density area (and provisioning them with 1 square mile PCS cellular areas), we need 2,438 MHz of spectrum at 4 bits/hertz and 2× for frequency reuse. This is still a lot of bandwidth. Furthermore, it requires substantial central switching/handoff capability, as well as low-cost base stations and premises equipment network interface units, so it is hardly “free.”
- **Switching—**If our goal is “dumb pipe/intelligent edge” capability featuring zero network switching costs, each end customer would need an optical or wireless “mega-set-top box” that met a $300 price point and could read a substantial fraction of the aforementioned 409-terabit datastream. This won’t happen soon.

Besides these supply-side limitations, there will be substantial demand-side growth to tax the system’s capabilities:

- There is large upside subscriber potential. Worldwide Internet subscriptions total about 80 million today—still a small number compared with 800 million total telephone users or 6 billion world population.
- There is substantial Internet usage demand upside. According to Pac Tel, the typical California ISP line supports 15 end customers and generates 20,000 minutes of use (MOUs) per month, or 1,333 MOUs per subscriber. This translates into a mere 3 percent usage on a weekly basis, compared with 4.4-hour-*per-day* TV usage. Therefore, Web usage levels have the opportunity to increase significantly as more user-friendly and higher-bandwidth sites are developed.
- There is substantial bandwidth upside demand. Currently, most Internet subscribers are limited to 28.8 or 56 kbps (excluding the odd ISDN line). If affordable broadband were available, however, many consumers would want service at T1 levels or higher.
- Substantially higher-bandwidth applications are in the offing long-term. For example, 1,080-line compressed HDTV would require four T1s; sending 32-bit uncompressed VGA computer graphics would require OC-12 transmission speed (402 T1s). Virtual reality applications under exploration at various laboratories require 1 Gbps of transmission capacity, 3D holographics require even more.

The final reason to be skeptical of supply-side answers is that data networking’s well-documented fractal demand pattern means that even if average capacity is adequate, there will be congestion during peak periods, hindering time-sensitive applications like voice.

Accordingly, we conclude that bandwidth will not be too cheap to meter, at least in the next couple of decades. Thus, we need to consider the MIT economist suggestions for congestion control.

Their approach would deal with the Tragedy of the Commons problem by charging for marginal usage. How practical is this solution?

Practicality of Real-Time Metering Solutions

We generally dislike real-time metering solutions. Proponents argue that metering (*a.*) is necessary to avoid gridlock congestion, and (*b.*) will not necessarily raise prices for the average user, since only marginal peak-load usage will be metered, while base fixed tariffs can be kept low.

These arguments may be correct in theory, but they won't work in practice. Given data networking's unpredictable traffic patterns, metering by time-of-day won't work, and therefore we would need real-time priority controls and pricing. Indeed, to make the system work, we would need to monitor each data packet for QOS priority (perhaps even each packet on each router hop) and price accordingly.

Unfortunately, besides having the technical QOS problems mentioned above, such monitoring likely is cost-prohibitive. Given that a data packet contains an average of 1,000 information bits and 100 control bits, a single OC-12 transmission line would require storing roughly 56,545 distinct 1,100 bit records each second, or over 60 megabytes per second.

An approximation of the dollars involved: an ordinary phone bill today costs about $2 per month to process. With 160 million phones, this comes to almost $4 billion/year. In fact, this is optimistic, since it ignores the development costs for billing software, and the revenue loss if billing software is not available for new services. A new packet-billing system would cost far more.

Even assuming that a tracking and billing system were perfected, we anticipate some thorny pricing issues, given that multiple hops often are needed to complete an Internet "call," and the Internet is not connection-oriented. Because of these factors, one does not know in advance how congested each hop is, nor do successive packets necessarily follow the same path as the first. If the system has a marginal pricing scheme to control usage at peak load periods, but a downstream hop has higher congestion than the origin hop, how will the system detect this in real time to restrict demand at origin points? In a connectionless system, this seems almost a contradiction in terms.

If we can't remake the Internet into a quasi-engineered network, we are left with the problem of how to price multihop calls. How should someone be charged for a 10-hop call if they pay for high-quality service and receive it in hops 1–5 and 7–10, but hop 6 is hopelessly overloaded, and inadequate QOS is provided? Will the customer get a rebate because of hop 6? Will the rebate be for the entire call, or $\frac{1}{10}$ of the call?

Beyond this, we shudder to think about how these packets would be summarized on a customer's bill. By packet, or even worse, by packet-hop? Such a bill could only be delivered electronically—each one would require thousands of pages of paper.

Alternatively, carriers could bill in aggregate. But if they did, how would customers audit and challenge the bill? If a customer sends information across multiple networks, each network requires interconnection agreements with one another as well as audit trails to demonstrate QOS, and this in turn implies significant electronic storage.

These issues raise fraud and quality-of-service concerns, such as:

- How will customers protect themselves against incorrect peak-period billing (since the new peak traffic periods will occur at semi-random intervals determined in real time by routing protocols, rather than the traditional telephony night-time and weekend discount periods)?
- How will customers protect themselves against the network provider unilaterally downgrading QOS level but charging for the higher price?
- Since most calls will pass through multiple networks, how will originating carriers be protected against a downstream interconnector unilaterally downgrading QOS but charging the higher price?

Another Approach: Class of Service

As people come to terms with the difficulties of real-time QOS metering, they have gravitated to a related approach called Class of Service (COS), supported by new protocols such as MPOA (for ATM) and RSVP (for IPv4). COS doesn't have a single routing queue where priority is determined based on real-time congestion and each packet's QOS header. Instead, COS establishes separate priority "classes," each with its own queue. Proponents often draw analogies to the airlines' first/business/coach class service.

COS is simpler to implement than QOS—it's not necessary to do real-time metering and billing. In the end, however, it is inadequate because:

- Multiple priority-based queues may help solve congestion problems *on average*, but they cannot deal successfully with *peak-load surges*. For any service class, there still is a single queue that prioritizes entries, and at any one time the highest-priority packet is being processed. If the fraction of time the link is busy grows too high, long queues and delays are inevitable.

To illustrate this by returning to the airline analogy: An airline with 100 seats overall might allocate 16 to first class. This prevents sale of these seats to mass market customers and increases the odds that there will be seats available at the last minute to high-end customers. However, it does not *ensure* first-class seating. If 18 people want a first-class seat, two will be out of luck, even if they are willing to pay the premium price. Accordingly, in telecom networks, priority service queues may be good on average, however, at peak times (for example, when the stock market nosedives), they will be overloaded.

Besides technical problems, real-time per-packet billing would be cost-prohibitive

FIGURE 2 Current ILEC Narrowband "Unbalanced" Architecture

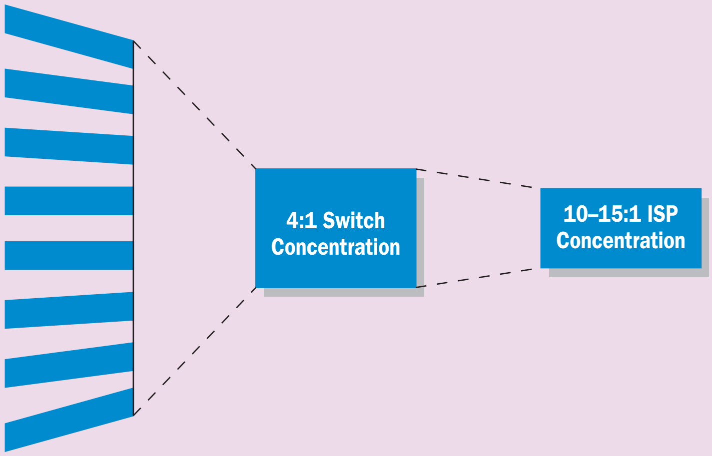

FIGURE 3 The Balanced Pipe Alternative

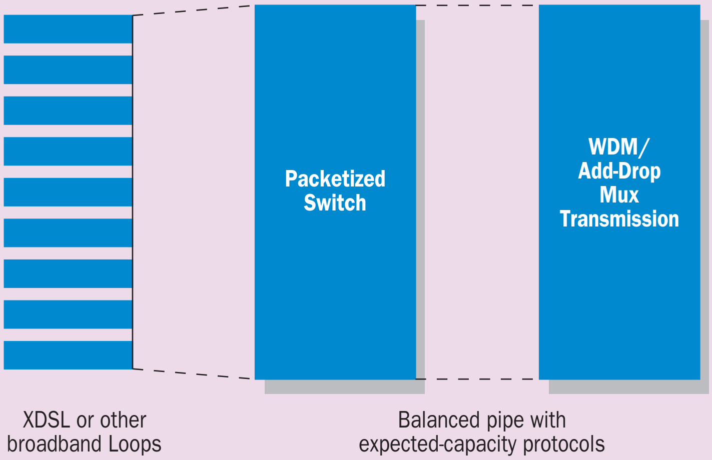

■ The same issues exist for COS as for QOS with regard to fraud in multihop systems.

In sum, COS is better than QOS, but it still does not assure Fedex-like "absolutely positively there by 10 a.m." service. If we want real-time voice service with decent quality, we need something better.

Why Advanced Network Architectures Are Needed

To deal with the need for peak-load throughput allocation and the impracticality of real-time per-packet QOS and COS metering, we propose a fixed price for a fixed guaranteed level of bandwidth. This option derives from a plan proposed by David Clark of the MIT Laboratory for Computer Science in the book *Internet Economics*, but we add elements based on our analysis of network economics. We advocate:

1. Engineering a network using advanced photonic and packet-switching technology for system-wide bandwidth capable of delivering a high level of service to all customers simultaneously at an attractive price. We call this the "Balanced Pipe" approach.

2. Using edge-shaping protocols to determine whether the network is running with much spare capacity. If so, allow everyone unlimited access to the network. If not, limit everyone to their minimum "expected capacity."

3. Because any Balanced Pipe approach can be sabotaged by interconnection with underengineered networks, isolate the above system in an extranet separate from the public Internet. Only allow interconnection with extranets employing similar Balanced Pipe/Expected Minimum engineering and protocols.

The Balanced-Pipe Network

The current ILEC network architecture is an unbalanced pipe, in which cumulative local loop capacity is substantially larger than switching and transport capacity (see Figure 2). The current dial-in Internet structure is even more unbalanced, with 10 to 15 customers per modem line. This structure does not handle congestion well. Overcapacity results in call blockage or, for packetized systems, call delays and packet loss.

Having rejected switch-based priority metering and billing systems, our suggested alternative is to create a pipe without constrictions—the "Balanced Pipe." This alternative involves fundamental changes in *switching, local loop* and *transport* (see Figure 3).

■ **Switching—**In the most basic Balanced Pipe configuration, if we have a LEC with 10,000 narrowband (64 kbps) phone lines, we would provision a non-blocking central office switch with 640,000 kbps of capacity.

For a network engineer with a "Bell-shaped head," this is economic anathema. Today's average RBOC gross switching plant costs $360 per line, $200 net plant, even using 5:1 line-to-trunk side concentration. Replacing this with a 1:1 non-blocking system would result in a gross investment of $2,000 per line and $1,000 net plant. Our hypothetical 10,000-line network would cost $20 million gross.

Fortunately, packetized switching offers a compelling economic alternative. To support 10,000 64-kbps phone lines, a packetized switch needs a capacity of 640,000 kbps. At equipment costs of $75 per T1 for a plain vanilla frame relay switch, plus $200 per T1 for a frame relay assembler/disassembler, the gross investment requirement is only *$117,000, or $11.70 per phone line.* This is a 97 percent cost savings over a circuit-switched network.

Indeed, it is even better, because the packetized construct is totally non-blocking. In our Balanced Pipe system, customers could leave their phone "on" 24 hours a day. This also resolves the infamous "ESP Exemption" issue that RBOCs complain about. With circuit switching, RBOCs argue that they are undercompensated for high-volume ISP lines. With Balanced Pipes, however, there is no marginal usage cost, so we don't care if they run 24 hours a day, seven days a week.

Indeed, packetized switch economics are so compelling that we can go one step further. Let's say that we want to provide each of our 10,000 customers with non-blocking T1, with 8 kbps dedicated for narrowband voice at an 8-kb vocoder rate. The capital cost for this is *$761,000, or $76.10 gross plant per T1 line:*

- 10,000 × $75/T1 for frame relay = $750,000
- 10,000 × 8 kbps/line 1500 kb/T1 × $200 per T1 frame relay assembler/disassembler = $10,667
- Total = $760,667

Thus, in our Balanced Pipe architecture, we can provide each customer line with T1 24-hour "on" capability, for around 80 percent lower cost than the gross plant of the existing circuit switched network.

- **Local Loop—**According to the FCC, current RBOC local loop investment per line is approximately $815 gross, $400 net. For the most part, this loop is non-blocking, and therefore supports our Balanced Pipe concept (either in its current configuration or upgraded to DSL). If/when someone installs fiber to the home/office, this too will be consistent with Balanced Pipe architecture.

In contrast, cable modems do not provide balanced pipe local loop, since the limited amount of dedicated data bandwidth, divided by the number of subscribers, could lead to congestion.

- **Transport and Muxing—**In the current hierarchical circuit-switched architecture, ILECs spend $355 in gross plant per wire for transport. As we have already noted, the system is replete with capacity constraints. However, with WDM fiber and advances in muxing and software, it should be possible to provision transport that provides an end-to-end Balanced Pipe.

Edge Shaping and Expected Minimum Protocols

In theory, a Balanced Pipe network will work without congestion. In practice, however, we need congestion control, for two reasons:

- **Recurring Behavior**: Let's assume customers have T1 lines, and a Balanced Pipe network is provisioned against that demand level. If customers begin to install DSL modems transmitting at 8 Mbps, the network becomes unbalanced.

In other words, customer demand for increasing bandwidth leads inexorably to "Tragedy of the Commons" behavior. We need a mechanism that prevents customers from overusing the network without permission.

- **Noneconomic Pipe Overbuilding**: If we build a Balanced Pipe sufficient to cover everyone's peak load demand, it probably will be underused most of the time. This, in turn, adds unnecessary system cost.

Accordingly, we need to control usage, to ensure that our Balanced Pipe network is not overengineered. In contrast to the IPv6 per-packet QOS header approach, our approach uses *edge shaping,* in which end customers are each limited to the bandwidth they have arranged to receive.

There are two ways to do edge shaping: with customer-located boxes, such as those developed by Xedia (www.xedia.com) and Packeteer (www.packeteer.com), or with originating switch software, such as that developed by Ipsilon (www.ipsilon.com) and NeoNetworks (www.neonetworks.com). Either way, customers purchase a specific amount of bandwidth that they cannot exceed without approval. If someone needs more bandwidth, he or she can contract with the network administrator to buy more.

Having thus controlled overuse, we deal with overcapacity by adopting David Clark's *expected minimum* protocol. In any data network, there will be frequent occasions of substantial system overcapacity, at which times there is zero marginal usage cost. Accordingly, per the expected minimum protocol, we are able to allow unlimited bandwidth use whenever the incoming switch is operating below capacity. This accomplishes several objectives:

- Each customer can decide how much priority bandwidth he or she want to purchase (that "absolutely positively" gets through even when system utilization is high). A budget-conscious customer might set this as low as the bandwidth for a single narrowband voice call.
- Each customer can make his or her own tradeoff with respect to maximum throughput and minimum available capacity. A customer with a T1 expected minimum but 8-Mbps ADSL capacity can test how well/how often he or she is able to transmit at 8 Mbps. If fast transmission is available most of the time, he or she may not feel the need to purchase an 8-Mbps pipe. Indeed, if the system is running low enough most of the time, an 8-kbps to 64-kbps line may be sufficient to ensure single-line phone service at all times, with data transmissions handled on an "as available" basis. On the other hand, if the system runs at high average utilization, more customers may need to buy 8-Mbps lines.
- Network capacity planning becomes relatively easy. We no longer have to worry about Erlangs and busy-hour CCS. We simply build out enough Balanced Pipe to cover the contracted-for expected minimum demand. If customers feel constrained, they can order more expected minimum capacity. If they want to save money, they can order less.
- With Balanced Pipes, customers can stay "on" the network 24 hours a day, seven days a week. The interminable arguments about "excessive" ISP modem line use become a thing of the past.

One important implication of the Balanced Pipe is abandonment of a key Internet tenet—the belief in nonengineered networks. Without central engineering, Balanced Pipes can't exist.

With Balanced Pipes, customers can stay "on" the network 24 hours a day, seven days a week

Private Extranets, Not Public Internets

Abandoning the Internet's faith in nonengineered networks is essential because the Internet, as currently structured, cannot provide a Balanced Pipe. Even if Backbone Provider A has a Balanced Pipe, other backbone providers without such pipes will sense that there is spare capacity, and route traffic to Provider A—making it, too, an unbalanced pipe subject to uncontrolled congestion.

What if we limit Internet peering to backbone providers with substantial capacity with no charge for OC-3 peering, but some charge for lower bitrate interconnectors? This is what UUNet advocates, limiting peering to companies with OC-3 backbones. But setting a particular level such as OC-3 has inherent problems:

- It is insufficient, because utilization, not absolute transport rates, is what's crucial. Two backbone providers may have equal transport speeds but grossly different capacity utilizations and network-to-customer demand levels.
- It is unstable, because even if we have two Balanced Pipes interconnecting on day one, there is a built-in incentive to underexpand one's own network and send traffic to the peering provider.

Accordingly, our solution does away with the public Internet as a high-quality communications vehicle (we would keep it as a low-quality vehicle of last resort). We propose sending traffic over a private extranet that does not peer with other companies on a UUNet model but might peer on the basis of each side buying ports on the others' networks, each operator including this bitstream requirement as part of his system's Balanced Pipe capacity requirement. This would preserve interoperability while preventing would-be parasites from unbalancing the pipes.

Conclusion

In the end, we don't believe in the practicality of the "too cheap to meter" concept, because we won't see infinite fiber capacity and 400-terabit home set-top boxes for $200 anytime soon. However, if we move toward a pricing-per-bandwidth port paradigm, we can avoid both extreme congestion on the one hand and complex metering solutions on the other. Of course, incumbent providers might find our suggestions threatening:

- The Netizen community has become accustomed to "Tragedy of the Commons" pricing and peering. It considers engineered solutions to be at odds with the idea of a "democratic" connectionless Internet.
- Telcos (ILECs and IXCs) make most of their current profits from toll charges. Therefore, it is understandable that they strongly prefer metered pricing for priority traffic.

However, we believe that free market competition will result in an inexorable industry shift to our flat-rated solution. Attempts to do otherwise are rearguard actions at best□