Straight Talk About Bandwidth Capacity

Net Forecasts – Peter J. Sevcik
BCR Volume 31, Number 9
September 2001

Bandwidth is the most talked about yet least understood aspect of telecommunications. All too often, investments have been based on misinformation and misleading advice about the quantity and form of bandwidth that is needed in the future.

So, in a nutshell, here's where things stand: There is a *temporary* glut of switched bandwidth in the Internet core, but contrary to the current conventional wisdom, there is no fiber glut (see this column in *BCR*, January 2001, pp. 10-11).

It's true that there are untapped reserves of long-haul conduit and fiber, but does anyone refer to the Middle East or Alaska oil reserves as a "glut"? In both the oil and telecom businesses, it takes a tremendous amount of work and investment to transform the reserves into a useful commodity. Moreover, the work can't be performed fast enough to accommodate a sudden change in demand.

For example, fiber is transformed into a useful commodity by lighting it with wave division multiplexing (WDM), and by adding switching equipment to create bandwidth a user can call upon as needed. But the same recent dramatic changes in transmission efficiencies that have spawned alternative views of how bandwidth will be provisioned and delivered, have also sparked debate about how to apply the capability of WDM wavelengths within a strand of fiber.

An Analogy That Doesn't Fit

One view is that bandwidth is becoming such an inexpensive commodity, that it's "OK" to waste it in order to save on another more expensive part of the architecture. Bret Swanson of Gilder Publishing made a strong argument for this position in a recent article in *The Wall Street Journal (see* "Bad Bets on Bandwidth," June 19, 2001, page A22).

Swanson used the analogy of the computer chip to what is occurring with WDM. Back when computers were expensive, they were shared among users; today, they are inexpensive and are dedicated to each user. Indeed, as a result of Moore's Law, we're seeing chips going into almost every man-made item. Mr. Swanson reasons that with bandwidth becoming inexpensive, someday, networks will supply dedicated bandwidth to each user and, inevitably, to each man-made item.

That analogy, however, is flawed. The microchip ushered in an age in which computing bandwidth (mega instructions per second, MIPS) became irrelevant, and innovation focused on ubiquity. Everyone owns many chips, most of which are hidden in products. No one cares how many MIPS of computing bandwidth each chip has.

Analogies are powerful tools, but if misapplied, they can be the source of a powerful mistake. The appropriate analogy between the chip and network revolutions is that innovation will give users and then "things" ubiquitous access to the Internet, and bandwidth will not matter. Users will no longer care or even know the data rate at which they are connected. They will only care that they are connected. Therefore, the appropriate analogy between chip and network technologies applies to *access* not bandwidth.

Shared vs. Dedicated Bandwidth

The misapplication of the chip analogy leads to a false set of assumptions:

1. Users will have an inexpensive, dedicated connection between themselves and the destination to which they're communicating.
2. This connection will need a fixed, modest amount of bandwidth that is sufficient to the task at hand.
3. Therefore, millions of connections will be needed.
4. And therefore networks should be built with many circuits that can be switched into place as needed. (Circuit switches establish dedicated bandwidth, while packet switches manage shared bandwidth.)

The circuit-switching (connections) versus packet-switching (datagrams) debate is not new (see, for example, "Why Circuit Switching is Doomed," in *BCR*, September 1997, pp. 33-35). Circuit-switching proponents have proposed the telephone switch, X.25 switching and ATM switching, while the packet-switching proponents converged on IP as the *de facto,* international, packet-switching standard.

But the old telephone switch network is migrating to an IP infrastructure, and X.25 and ATM are falling by the wayside. Users and devices *will* connect to the Internet with packets not circuits, obviating the need to provide millions of circuits to each user connection. It is already too late and impractical to provide a dedicated connection to a user session.

For example, a typical Web page loads by providing information from about a dozen sources. Packet switching provides access to the many sources without the overhead of a connection setup. Each source supplies its data within fractions of a second using a shared bandwidth pool.

The other reason to build a long-haul network with many small circuits is to support the enterprise private line market. However, that market is already transitioning – first, to the shared model of frame relay, and next to shared VPNs. The private-line market will always migrate to the lowest cost approach, there's been a progressive acceleration of cost declines by SONET, then ATM and finally Ethernet networks. Ethernet is the clear long-term, low-cost winner in transmission precisely because it is packet-based and takes advantage of shared bandwidth pools (see this column in *BCR*, September 2000, pp. 12-14),

Maintaining High Spectral Efficiency

A fixed-sized guard band of unused bandwidth must isolate each trunk or wavelength within a WDM spectrum. The approach of supplying many small trunks therefore requires many guard bands bandwidth between each trunk (low spectral efficiency), while the approach of a few large trunks reduces the number of guard bands (high spectral efficiency). In short, providing many small bandwidth circuits as part of a circuit-centric strategy leads to low spectral efficiency.

So, users will use packets to connect into a shared metropolitan bandwidth pool and then be aggregated for connectivity into the long-haul carriers. The long-haulers will receive traffic in ever-larger aggregations of users and bandwidth load, and these large access pipes must in their own right be trunked onto yet larger long-distance trunks.

That's going to require a bandwidth hierarchy, which, as you move from user through access and into the core builds into ever-larger bandwidth trunks. The flexibility of the long-haul carriers will be determined by their ability to create, deploy and re-deploy these high-capacity trunks out of the few strands of fiber that will be lit each year.

These trunks may use some form of virtual circuit to assign a wavelength to a traffic bundle between major geographic regions. The emerging method for this form of traffic engineering is MPLS and its *lambda* derivatives. These very high-capacity virtual trunks are taking bandwidth out of a shared pool of bandwidth available in the lit fiber plant.

High spectral efficiency systems will allow carriers to economically build large pools of bandwidth which can be filled by packet switches. By contrast, low spectral efficiency systems will put a carrier at a significant economic disadvantage, and force carriers to mine more of the fiber reserve to provide an equivalent amount of bandwidth in use. The only bandwidth that generates revenue is bandwidth that's used.

Just as oil companies have found ways to get more gallons of gasoline and other products from each barrel of oil, the carriers must get as much bandwidth in use out of a fiber strand. That means deploying large bandwidth pools that the packet switches can fill. Dedicating bandwidth to a single user or even a single access network would be extremely inefficient. Carriers do not need thousands of unsold small pipes.

Keeping Capacity Ahead of Demand

Carriers need to be able to put the right amount of capacity where and when it's needed, and to accommodate unexpected traffic spikes and shifts in demand. The primary customer of the long-haul carriers will be large companies and metropolitan carriers that aggregate traffic. The size of the customer circuit will continue to be larger. The bursty nature of the packet traffic on these large access circuits will magnify the shifting demands of new applications on the Internet.

With access users contending for each trunk, the average trunk must operate at speeds significantly higher than the highest access line rates. With today's edge technology interfacing to the core at up to 10 Gbps, carriers need trunks that go 40 Gbps as a starting point to keep up with the growth of access bandwidth. Carriers that fail to do this will not only experience degraded network performance, they run the risk of having insufficient capacity to upgrade existing customers or provision new ones. Not only will this stifle growth, it will prompt customers to switch service providers. (While there's a niche market for hundreds of 1-Gbps wavelengths, it only applies to the "last mile" – terminating many business customers in a single location, e.g., a skyscraper in Manhattan.)

I Want Capacity, Not Wavelengths

Customers want to buy capacity from carriers, and they want the flexibility to buy in the increments they need. Unfortunately, the existing optical network systems lock carriers into selling fixed – not flexible – capacity wavelengths, because in the past a carrier could sell only an entire wavelength.

But this limitation is going away. Advanced optical technologies will allow carriers to sell capacity in true bandwidth increments rather than by the wavelength.

The approach of selling many small wavelengths leads to poor spectral efficiency. It not only puts the carrier at an economic disadvantage, it also puts the carrier in the most awkward of all market positions: selling the wrong service relative to demand.

Peter Sevcik is president of NetForecast in Andover, MA, and is a leading authority on Internet traffic, performance and technology. He has contributed to the design of more than 100 networks, and led the project that divided the Arpanet into multiple networks in 1984, which was the beginning of today's Internet. He can be reached at peter@netforecast.com.

NetForecast measures application profiles, user behavior and network statistics to predict performance, adoption, and market impact of new technologies. The firm has helped leading service providers, enterprises, and vendors navigate the changing competitive landscape of the Internet economy. We supply key technical and market guidance to ensure the success of network-based projects, products and services.

www.netforecast.com