A Contender With Staying Power

Net Forecasts – Peter J. Sevcik
BCR Volume 30, Number 9
September 2000

Three years ago, I wrote an article that contained a cost/performance analysis of four leading switch technologies over time (see "Why Circuit Switching is Doomed," see *BCR*, Sep 1997, pp. 33-35). The fundamental conclusion of that article was that packet-based solutions would eventually dominate, because the cost of switching a bit per second doubles at a faster rate than circuit switches.

Some readers of that article assumed that the operational costs of a network would reflect those same curves, but they forgot one key point: The cost of a switch is only a small portion of the overall cost to operate a network service.

Figure 1 presents the total monthly cost of a Mbps of bandwidth sold, over time, on networks built from alternative Layer 2 technologies -- SONET, ATM or the emerging Gigabit-Ethernet (GigE) option. A SONET network provides dedicated circuits, while ATM and GigE networks can provide dynamic connectivity across their respective networks. The data in the figure is based on models for full-duplex bandwidth, including last-mile access to an office building in major metropolitan areas.

The data reflect the total costs to a network service provider (NSP) that are required to operate a self-contained, national business-to-business data service. Each model represents a complete cost structure including research, deployment, operations, maintenance, billing, customer care, leases and, of course, capitalized hardware including switches. Note, however, that the models do *not* present the price buyers would actually pay for these services, they do not include operating margin or profit.

Each of the modeled networks provides full service up to the level of bandwidth purchased (e.g., buying 10 Mbps on a 45 Mbps access line); there is sufficient bandwidth within each network to support a high level of quality and availability. The varying ability of each technology to provide a multiplexing gain on the trunks also was taken into account – i.e., protocol and switching efficiency.

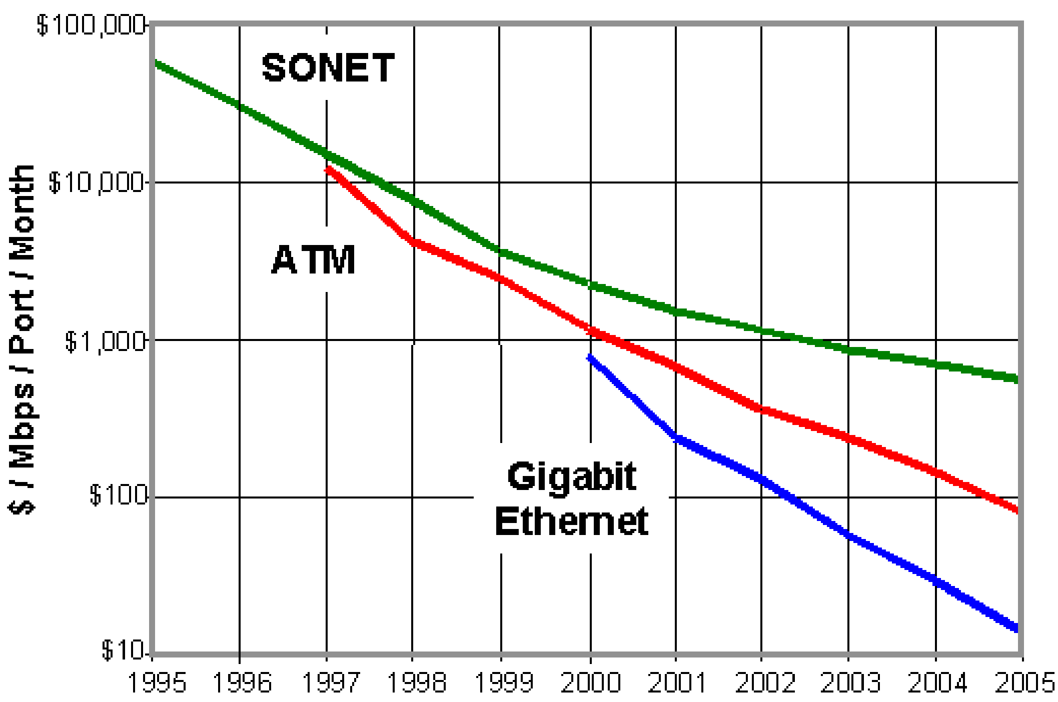

Figure 1 – Total Cost to Deliver Bandwidth for a US Network Service Provider

If telephone calls are converted into data, these costs can be used for trunking VOIP traffic. But the curves do *not* include any telephone service features, conversion to VOIP or the costs to gateway from the PSTN into a pure data service.

In short, these models depict what it takes for an NSP to serve high-volume business customers – customers whose requirements mandate the availability of an *under-subscribed* service. The data cannot be compared to the over-subscribed services of ISPs – e.g., consumer DLS or cable access. Moreover, the models reflect the cost structures of new, "greenfield" service providers, and therefore provide a way to compare traditional carriers/service providers, the CLECs of the mid-1990s and those who are emerging in the early 2000s.

Observations from the Curves

While each of the curves in Figure 1 could be analyzed in terms of different scenarios, the most critical factor is the target market for the network: How fast it buys ports and then how fast it buys bandwidth.

When I looked at Figure 1, I found that the slopes of the curves were not as dramatically different as the slopes of the underlying technologies I published three years ago. The impact of declining technologies, even one that is declining rapidly, is buffeted by all the other operational costs.

In absolute terms over the next five years, GigE delivers bandwidth at a cost generally 70 percent cheaper than ATM, and ATM is itself 70 percent cheaper than SONET. Since the technologies appear to maintain a similar pattern as they try to overtake the older competitor, the most important factor is how long each technology can sustain the race. GigE has the edge, because it cuts the cost per bit by 50 percent every 10 months, which is significantly faster than its rivals.

Installed Base – Anchor or Opportunity?

After performing a sensitivity analysis on the various network models, I found that the key issue is the ability to sell more bandwidth. Since an NSP's network often has more capacity than can be effectively delivered due to the limits of the access interface, in order for NSPs to sell more bandwidth they have to convince customers to upgrade their existing interfaces or even complete services.

Now here's where things start to get interesting. While every data service must provide high availability and low error rates, it's important to remember that data services are featureless (no call waiting, etc). However, in order to sell the only product that has true value – bandwidth – NSPs force the customer to change CPE interfaces (e.g., modems on analog lines, then dedicated digital circuits like T1/E1, X.25, Frame Relay, ISDN, ATM and now GigE) in order to attain higher levels of bandwidth.

By contrast, telephone customers do not buy more bandwidth, phones have required 2 kHz of bandwidth since the days of Alexander Graham Bell, but telephone customers tend to need many more features. Data customers, on the other hand, buy more bandwidth but no features (a reliable bit transfer is still a reliable bit transfer).

This explains why voice and data networks will always be treated differently: The same CPE instrument you bought in 1970 is still working on the telephone network, while during that same period of time, you've had to change the data CPE instrument many times over.

The data services listed above all have a bandwidth dynamic range – i.e., the slowest to fastest data rate the service can support over its lifetime -- of 16:1 to 24:1. However, there are likely to be speed bumps along the way requiring the purchase and installation of new CPE as the customer moves up the bandwidth hierarchy. For example, each speed increase in the analog modem world (2,400 to 56,000 bps) required a new modem, and in order to move up the SONET hierarchy's 16:1 dynamic bandwidth range (155 to 2,480 Mbps), customers have to buy a different – and costly -- interface at each step (OC3, OC12, OC48).

But the marketplace reality is that data customers do not change fast enough for NSPs to race down the price/performance curves as fast as they'd like. The NSPs have to keep operating older, slower services and technologies as long as its installed base of customers want to use it.

So there's the trap: The per-bit cost of operating at slower speeds is higher than operating at higher speeds. An installed base of customers keeps the NSP's inherent average costs high.

The data refutes the conventional wisdom that a large, older NSP enjoys inherent economies of scale. Instead, because all the "legacy" data services are still in wide use on a global basis, the established NSPs are cursed; they're simply older service providers with older customers.

But past is not necessarily prologue: New NSPs who base their service solely on GigE can avoid this trap. GigE is the first WAN technology that can support a dynamic range of 1000:1 with only one speed bump (at 100 Mbps). It will be much easier for these NSPs to sell more bandwidth to their installed base than any of their predecessors, and that means that instead of acting as an anchor, the installed base will become an opportunity to sell more of what customers actually want: bandwidth. The broader bandwidth dynamic range with the ability to attract a higher bandwidth-growth customer is the most significant factor separating the curves in Figure 1.

Coping with the Real Rate of Technology Change

To help put this in context, new data switching technologies are doubling in performance/cost every 20 to 10 months, a pace of improvement that is faster than the computer chip. In order for a business to maintain a competitive edge with its information technology, it needs to upgrade computers every 36 months, because, per Moore's Law, computers are doubling in performance every 18 months. Any IT department that really takes this concept to heart would therefore replace one-third of the computers every 12 months. Every year one-third of the computers are thrown out. Gone.

Applying this same principle to the NSPs, they should also be retiring at least one-third of their equipment every year. None do. Instead, older equipment moves closer to the edge while the new, higher-performance equipment takes over at the network core. The net result: Old costs aren't written off, new technology – and associated new costs -- are added; the number of network layers and the associated complexity increase.

All of which puts the older NSPs in a real dilemma. Their installed base isn't motivated to change to newer services, and it's hard to imagine them announcing that older customers must convert to Ethernet-based service within three years because, at that time, all older forms of access will be terminated. If the NSP is lucky, the business climate will stay good and they'll be able to sell their older customer base and service to another NSP.

Again, the Ethernet-based services are at an advantage. They can start a customer at 1 Mbps and then increase the sale of bandwidth to that customer -- doubling it 10 times before running out of a service to sell. Even if demand for bandwidth from business customers were to double every 20 months, the NSPs get 17 years to sell more service while the previous technologies run out in about 7 years.

More Cost Savings Required

But bandwidth alone isn't enough; to attract customers to the new network, NSPs must make the price attractive. Despite the technical advances we've all enjoyed, the cost of a Mbps is still not dropping fast enough to open the floodgates to new bandwidth-hungry applications.

All of the NSPs -- even those that base the networks on GigE -- will have to find new ways to hold costs down. This is why techniques like self provisioning -- the customer buys more bandwidth via a simple Web interface – have generated so much interest. In the not too-distant future, customers in office buildings will be able to simply plug into a patch panel in the telecom room to become a new customer of a service. No technician will have to visit the client, and building landlords or a business organization might serve as "agents," selling on behalf of NSPs. There'll be more promotions, group buying plans and bundling as service providers seek to gain new customers who'll bring new traffic onto their networks.

All of this, however, is predicated on the continuing growth of bandwidth required by business customers. Current trends indicate that GigE NSPs have a long period to gain market share without becoming "incumbents" who are burdened with a pesky, old, installed base.

However, from a GigE perspective, even if the growth levels off, the outcome doesn't change. If the growth curves flatten, there will be no need to develop yet another "next-generation" access method. The public data network evolution will end with Ethernet as the last entrant, but the only one left standing.

Peter Sevcik is president of NetForecast in Waltham, 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