

Service Providers: Give Customers What They Want, Not What You Want

Net Forecasts – Peter J. Sevcik

BCR Volume 31, Number 7

July 2001

Service providers and their equipment vendors are continually planning the next new service or variation on an existing service, but I am often struck by how consistently they misunderstand the buyer-seller relationship. Network services are purchased in a food chain of buyer-seller relationships; a broken promise is a broken link that hurts the rest of the chain.

The Need For Flexible Bandwidth

Let's start with a simple notion: Network service providers (NSPs), first and foremost, sell bandwidth. It is sold in increments of at least two-fold (buy another equivalent circuit to the one you already lease), four-fold (buy the next OCx in the numbers progression) or as much as thirty-fold (buy a DS3 to replace a DS1).

The problem is that these extremely large increments do not match the growth of demand as it's experienced by most customers. Most enterprises, not to mention regional carriers and smaller ISPs, need bandwidth access in much finer increments.

The math is simple: If you accept that traffic on the Internet is doubling every 11 months (see this column in BCR, January 2001, pp. 10-11) then you also must accept that traffic demand is generally growing at the same rate in all aspects in the Internet, including the traffic growth of customers. Therefore, the increments in which bandwidth is sold only match customer needs once every one, two or five years! Not a happy prospect for either NSPs or customers.

More granular bandwidth offerings, best exemplified by frame relay, don't apply to Internet access. An Internet customer needs variable bandwidth access to the entire Internet, rather than to specific locations as defined by a CIR on a permanent virtual circuit.

An NSP that sold bandwidth in flexible increments which reflect the real monthly shifts in customer bandwidth demand could gain a dramatic overall revenue advantage. Traffic models we've developed show that if two NSPs started selling

services on the same day, the carrier selling in smaller increments would achieve a 10-fold bandwidth sales lead by the end of the first year; even in the fifth year, the bandwidth sales lead would be 55 percent. In short, today's service providers are leaving money on the table.

A Chronic Mismatch Of Bandwidth

Given this sorry state, customers have little choice but to either over- or under-provision their Internet access. While there are short periods of time when the access bandwidth actually matches the customer's demand (kind of like a broken clock showing the "right" time twice a day!), most buyers are under-provisioned most of the time.

This fuels the following chronic condition: Local networks (LANs or MANs) are often over-provisioned, because they are built with capital budgets rather than re-occurring operating expense budgets. However, when LAN traffic hits the under-provisioned WAN, trouble is always brewing. This well documented speed bump can be managed with QOS techniques, but usually users simply live with uneven performance.

No matter how it's provisioned—fixed-size or flexible pipes—there is always a specified amount of bandwidth between buyer and seller. What is not fixed, however, is the traffic carrying capacity of the NSP behind the access line.

NSPs are notorious for oversubscribing a bandwidth pool. Few customers can really fill the bandwidth they think they bought, because somewhere within the NSP's network or at a large up-link to an Internet backbone, there is not enough bandwidth to support all the traffic. So, like the airlines, NSPs over-sell or over-subscribe their service, betting they can get away with it most of the time. If they consistently under-subscribed the network, customers would receive great service but NSP profits would be a little smaller.

Again, there are many ways to cope with this problem—using sophisticated classes of service and many of the same QOS tools customers use. But more often, the NSPs fall back on fancy SLA

language, which only a few customers have the means to measure and enforce. The default mode of operation is as follows: Packets are lost when congestion occurs, and TCP sorts out the problem by slowing down all the users affected. You and I wait until the process, page or job is completed.

Biggest Market Has The Biggest Problem

Given the above, regular readers of this column must think I have done an about-face, because I constantly gripe about the over-built Internet core; there is way too much bandwidth relative to the demand. But while in the aggregate supply exceeds demand, many buyers can't get the bandwidth they need, and people buy bandwidth to serve a wide range of applications and user constituents. It turns out that very high-end and very low-end customers can usually get what they need, with relatively little trouble.

For example, folks working in small offices and individuals who work on their own have a reasonable understanding of their usage and needs. However, this group creates a very small load on the large pool of network capacity. Given proper planning, most of this group can buy enough Internet capacity to satisfy their needs, particularly if they're located in metropolitan areas of the U.S. where there are alternatives from which to choose for Internet access.

Meanwhile, customers buying on behalf of a million or more users have access to serious capacity planning capability, and they can get the attention of many NSPs. In fact, these buyers can do what the NSPs do—play the statistical game of oversubscribing, because the “law of large numbers” often will work in his/her favor.

By contrast, the customers in the “middle,” get squeezed; if you're buying on behalf of a single enterprise or if you're a local ISP that needs to support a few thousand users, you've got trouble. You're part of a larger group that someone else has aggregated to access the Internet. You've got a user population that's too small for leveraging bandwidth economies of scale, but large enough to have a highly variable traffic profile. Your users are likely to experience congestion periodically nearly every day.

This is the most prevalent customer group coming to the wholesale and backbone NSPs, and it's the largest market for backbone NSPs. But since there's no way to buy bandwidth dynamically and/or flexibly, this large group remains chronically under-provisioned and over-subscribed. Congestion is part of the daily routine, and if it were to suddenly disappear, both buyers and sellers would be investigating why—either the seller under-subscribed his/her assets or the buyer over-provisioned his/her resources. Either way, someone would be leaving money on the table.

And we also need to remember that buyers often have good reasons for under-provisioning a site, even though it places users and applications in jeopardy by permitting congestion to occur. A network manager's job is to convince management that it is acceptable, within the business context of the site, to use QOS to make the best of a tough situation.

For example, many multinational firms are setting up engineering groups in remote corners of the world, such as India and Russia, where talent is plentiful and relatively inexpensive. These sites may have 500 engineers toiling on critical designs but they'll be supported by a fraction of the bandwidth available at headquarters. QOS techniques are the only option to make sure critical traffic gets carried at the expense of the less critical.

By contrast, service providers/carriers only use QOS to try and cover up the fact that they're over-subscribing their service. Nothing infuriates customers more than when an NSP says: “Please pay extra for the QOS features in order for your users to actually get the service you thought you were already buying.” The service providers have tried to spin this and hope that they'll fool customers into paying more for the same.

It's important to recognize just who chose to put the users and applications at jeopardy by permitting congestion to occur— that decision is usually the buyer's choice, and only the buyer can take steps to get out of trouble. Service providers should not cloud the waters by over-subscribing and then adding QOS to “fix” their problem.

Conclusion

This dance will prove a waste of time. The NSPs are searching for services that will be considered “value-added,” and thus produce higher margins to offset the falling price of commodity bandwidth. Unfortunately, they’re looking in the wrong place: Coming up with variations on bandwidth itself won’t get them where they want to go.

For example, we learned with basic telephone service that the market will no longer support these simple variations. It wasn’t all that long ago that telcos had differentiated charges based on time of day, day of week and distance. Today, those pricing schemes are history. Why should the data world replicate a model that’s already out of date? It’s wishful thinking by NSPs to think they can recreate simple differentiated data services based on peak rate, traffic during network congestion and destination.


The long-haul NSPs need to get real, and stop trying to fix the problem with differentiated services and marketing gimmicks. Instead, they should give the access ISPs and enterprise

customers what they really want: Flexible bandwidth at a reasonable price. If the bandwidth matches the aggregate demand, you don’t need any QOS or fancy SLA language. There are value-added services, e.g., security and personalization, for which NSPs can charge additional money, but they need to be put into perspective: They’re the equivalent of voice-mail and caller ID.

It is time for NSPs to separate services into two piles: Those that look good to the accountants and those that look good to the customers. You know which pile will really make money.

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