

The Premium Service Trap

Net Forecasts – Peter J. Sevcik

BCR Volume 29, Number 3

March 1999

The industry is in danger of falling into a trap: Network Service Providers (NSPs) are trying to create a market -- premium service -- without having developed a convincing case that customers will really benefit.

To a certain extent, the NSPs are being pushed into this trap, because the conventional wisdom among financial analysts is no money can be made selling cheap, basic Internet service. Salvation, according to this world view, requires that NSPs sell *differentiated* services, higher quality -- and higher margin -- services to business customers.

That may sound good, but there are some basic unanswered questions: Will customers buy "premium" service? And will they be willing to pay a premium price?

There's certainly no lack of technology being developed. Lots of equipment manufacturers want to show why and how their latest breakthrough can provide the foundation for differentiated services. But the techniques are so complex, that I always wind up asking: "Which grouping of variables do you recommend for Service A vs. Service B, and can a user really tell the difference?"

Their inevitable reply: "Well, that is for the network service provider to determine. We only supply the tool by which thousands of different services can be defined."

Recently an NSP pitched me on a differentiated service story. I was told: "We have different classes of service, each of which is defined by a set of queuing, precedence and buffering schemes."

I thought to myself, "Great! Someone was really trying this stuff." But when I asked whether the *user* would be able to see or experience a difference, the best they could reply was: "We presume so."

Basic Can Be Better

There's reason to be skeptical. Several months ago I finished a detailed analysis of service classes on a live network running among four distant cities. I used application profiles to represent transaction processing, Web-based applications and file transfers.

The results were surprising: *In general, the basic service performed better than the premium service.* All the traffic traveled across several hops of a shared network. The premium service was pre-assigned to better paths while the basic service was left to fight over the rest of the resources. However, that basic "best effort" technique does a good job at finding the best path in real time. So, when the network is lightly loaded, the basic service actually wins over the premium service most of the time. However, if the network is congested, the premium service maintains its level of service while the basic service degrades.

Since basic service usually performs better than premium service, it looks like differentiated services will be sold as insurance against those times when the network is having a "bad day." So, the first question a prospect ought to ask is: How often does a bad day occur?

There are several ways an NSP salesperson would respond. First, he/she can assert that *the NSP* is not having a bad day, but that overall network traffic level is beyond its control; like the surge that occurred to read the Starr report. Sort of like blaming a delayed flight on the weather.

Second, the NSP rep would probably try to shift the blame, suggesting that it's the customer who is having a bad day, and that the cause is the customer's failure to manage its traffic requirements. Or that the "bad day" is the result of poor planning for client-induced traffic surges (e.g. end-of-month roll-ups or surge in trades because the stock market is falling).

Finally, there will be times when the NSP will have to admit that it has lost a major trunk or

experienced a system failure that reduced its ability to handle all the offered traffic. This is the main reason to invest in a premium offering -- as insurance against the inevitable times when an NSP's service isn't up to par.

Why Make a Sow's Ear out of a Silk Purse?

The predicament over premium service started with the notion that the Internet was a marginal service, used by the technical elite but not ready for prime time. It was full of hackers and would melt-down any day.

Well, the Internet was never that bad, and the statistics on average delay and packet loss rates are improving. In short, a funny thing happened on the way to differentiated services: The 'Net got really good.

For example, AT&T just announced a new Internet premium service, which promises a packet-loss rate of less than 1 percent and round-trip delay of less than 80 msec. It's interesting to note that on the day of the announcement (January 26, 1999) AT&T's *basic* service was providing a packet loss rate of 0.08 percent and a delay of 51 msec. So, AT&T finds itself asserting its new premium service will consistently perform better than the basic service, but there are no measurements to back up that claim. In fact, the economies of scale presented by the large population of basic users versus the small population of premium users suggests that basic users will do as well as premium users most of the time.

The overall quality of basic Internet service even holds for a demanding application like international voice-over-IP (VoIP). For example, VIP Calling, Inc. (Burlington, MA) provides toll-quality VoIP services to more than a dozen locations including Athens, Tel Aviv, Hong Kong, and Seoul using its Assured Quality Routing (AQR) technique. If a path can't support toll-quality VoIP, AQR moves the call to one of several traditional, international carriers that are pre-provisioned as back-up to the Internet. But today's Internet is so good that it carries more than 90 percent of the calls -- an impressive level of call-completion on thin international routes.

Still, the NSPs feel compelled to embark on the premium-service path despite the possible serious repercussions. Here are a few scenarios:

- The NSP sells premium service as insurance against disasters that would seriously impact mission-critical applications. This is an honorable strategy, but likely to be limited to a small market which will have a little impact on overall NSP profits.
- The NSP could try to motivate existing customers to move to the premium service by carefully reducing the performance of the basic service. This strategy carries the huge risk of losing the fundamental customer base upon which any premium can operate. The NSP may soon be out of business.
- The NSP can be so successful at developing a good marketing campaign that it builds a significantly larger pool of premium customers than basic service customers. At that point the NSP will likely change its image into a premium NSP and become a one-service company.
- Finally, the NSP may employ a lot of marketing spin for a very small service improvement. This strategy does not require any significant investment in differentiated services.

Over the next 5 years the differentiated services market will not change very much. There will be a few major players, each of whom will excel at a limited number of capabilities. Premium services will cost more, and the very definition of "premium" will vary depending on the respective strength(s) of the NSPs. Added value will be derived from co-branding, bundled services, affinity marketing, etc., and while those are nice, none make it possible to deliver differentiated services from a single network.

We should all understand that Internet transport will continue to be a commodity business. The data crowd likes to boast that some day, people will buy Web services with phone service thrown in for free. But since we now operate on Internet-time, how long will it be before people buy an entertainment service with both the Web and phone services thrown in for free?

Where Differentiated QOS Is Really Needed

Now, don't get me wrong -- QOS can be important, and it will continue to be needed in

localized, over-subscription situations such as access and multi-service gateway points. The most obvious application is high aggregation of customers onto an NSP service.

New access methods permit several levels of aggregation from the home or business to the NSP. There is a clear economic incentive to over-subscribe the NSP's up-link. An access provider can sell premium service over basic service by putting users into a less crowded pool. This is equivalent of operating a standby-only airline with two lists for passengers for a flight out of town. But once all the passengers get on the plane (the NSP), they all arrive together at the destination.

Another QOS access application involves businesses using a single line on which they send traffic to different NSPs -- e.g., data to an ISP and voice to the PSTN. There is value in carving up the over-subscribed bandwidth since it is carrying fundamentally differentiated services. This is equivalent to providing interstate highway exits to the airport but only basic street access to the bus terminal.

Note that all of these applications are localized within a single office or line. In each case, the differentiated services are governed by the basically different services coming to the location. Furthermore, assignments to the different services can be quite static. These are much more tractable problems than trying to provide end-to-end differentiated services across a dynamic wide area network.

The bottom line: Customers need to develop their own business case for premium service, where the cost is justified based upon business requirements and demonstrable, understood network conditions, and where the customer can prove that the service is actually better. Otherwise, if you just buy the marketing pitch, you may be falling into the premium service trap.

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.