

Performance-Based Network ROI Analysis

Net Forecasts – Peter J. Sevcik

BCR Volume 29, Number 7

July 1999

"The network is so slow it is costing us money. Fix it or I will change suppliers!"

Those are two of the most dreaded statements a network service supplier -- either commercial or an internal networking/IT shop -- can hear. Any reply to a disgruntled customer needs to begin with a thorough analysis of the situation, which in turn means you have to understand just how slow is too slow and how much money does it really cost the business.

There are also good reasons for a buyer of communications services to perform this type of analysis. While networks are upgraded to accommodate a new function or growth, most of the time, upgrades are justified on the basis of better performance. But how can you evaluate whether you should purchase the better service without tying it to your own business needs? Both buyers and sellers of advanced services need a business-driven, performance-based, return-on-investment (ROI) model.

Finding the Business Cost

Since the overwhelming majority of business applications are transactional, completion time is the fundamental driver of success. It can be daunting to try to define response-time service requirements of all the enterprise applications, but there are ways to triage the problem.

For example, applications and users can be divided into two groups -- production and non-production. Production users -- i.e., positions in manufacturing, logistics, back-office operations, telemarketing, call center and customer service -- must keep up with the daily, even minute-by-minute, demands of their job; the enterprise's customers are directly impacted by the slow applications. By contrast, non-production users can compensate for a slow system, and the enterprise's customers

never see or know that anything's happened. For example, slow email or a slow planning application or a printer causes inconveniences, but is not likely to create a make-or-break situation.

A second way to segment the problem is by the type of expense incurred. Was the expense internal -- where lost productivity leads to lower profit - or external, i.e., an expense that lowers the revenue of the business? Segmenting by type of expense focuses attention on three specific costs that are associated with poor application response time: additional support, lost productivity and lost revenue.

Additional Support: Sooner or later, no matter how low user expectations have been set, a threshold for system response time is hit and users will call the customer service center or help desk to complain. How much does it cost to service a complaint? Well, the answer varies, but members of The Help Desk Institute in Colorado Springs, CO (www.helpdeskinst.com) report that the average cost to handle a service request is \$20. Many outsource vendors charge \$25-\$50 for a support call. Organizations with no help desk and that rely on the end-user to figure things out probably pay the most in hidden costs.

Lost Productivity: If the load consistently overwhelms a network system, eventually, workflow becomes affected enough that it becomes necessary to hire additional staff just to keep up. A financial services firm in Boston had to add an extra shift of workers when the back-office operations were slowed by poor response time, even though it would have been much cheaper to fix the system.

Some operations know how much each second of response time costs. For example, a telemarketing organization knew that once a three-second trigger was

reached, each additional second of response time translated into a requirement for an additional active seat, and that's a significant cost.

Lost Revenue: Some users and applications really drive the revenue of the business. If the system is slow, customers go elsewhere, and transactions or sales are lost forever. This used to apply to airlines, banks, stockbrokers, catalog sales and any call-center application, but today, as 'Net-based revenue climbs, so does the exposure to lost revenue due to poor performance.

Indeed, this has become so important that Keynote Systems, Inc. (San Mateo, CA -- www.keynote.com) has been publishing the Keynote Business 40 Internet Performance Index for the last 2.5 years. Their most recent Web Broker Trading Index measures the average time required to buy or sell stock on popular on-line trading sites around the United States. In the first week of operation, the index showed the worst site is 5-times slower than the best site; dial-up users are likely to see the worst site as *more* than 5-times slower.

The implications for on-line trading firms are profound: Customers will move to the better-performing sites, assuming the rest of the user experience is equivalent and service is priced nearly the same. In a commodity business, service becomes the key differentiator.

Any serious on-line business will have to learn the response-time-trigger that will cause business to be lost and customers to move. These triggers can be derived by correlating business income and net performance on a historic basis, or the triggers can be learned through user surveys and demonstration tests.

From there, the calculations are well understood: The cost of a lost sale is simply the value of an average sale. The cost of a lost customer is the expense required to acquire a new customer, plus

Net Forecasts

the lost business for the time it took to replace the customer. These tend to be values that are related to the average churn and the typical transaction of the business. Losing a customer at Goldman, Sachs & Co., for example, is probably much more costly than a lost Amazon.com customer.

Calculating the Return on Investment

Once you have determined the trigger response-time threshold for each category, you take the number of users with performance slower than the trigger times and the number of incidents per month where such a condition exists, and multiply that value by the cost per incident; that gives you the business cost. The trigger values and the sum of the business costs provide a practical metric for use in shopping for ways to improve performance.

Any network improvement must first show that it improves performance enough to eliminate the trigger incidents. And that's a test many improvements will fail to meet. The investment cost for a good design is the amortized capital and system-integration expense required for the improved communications along with recurring expenses. This reduces both the investment and the business value to a set of comparable monthly costs.

Example 1: The first calculation is quite familiar. For example, let's say that a remote campus generates performance calls to the help desk that cost \$10,000 month (20 calls/day x 20 business days x \$25/call). An upgraded network connection will cost \$8,000 per month (\$72,000 capital/36 months plus \$6,000 recurring monthly costs). The return is \$2,000 on an expense of \$6,000, so the ROI is 25 percent.

Example 2: The investment may be viewed as insurance -- it protects against critical periods that cause response times to shoot past the triggers. For example, a daily production order is sent every morning to a factory management system that must have all the orders to calculate

the detailed workflow. Each order takes 5 seconds to transfer.

On normal days, this system works just fine, but there are also days when network congestion adds 10 seconds to each order. Since there are typically 1,000 units, this translates to a 2.8-hour delay of production. It cost the company \$10,000 per hour to operate the factory, so each incident costs nearly \$30,000. When such incidents climbed to about one per month, an upgrade that cost \$20,000 per month was considered a good investment with a 50 percent ROI.

Example 3: A shipping firm has to get its ships containing highly perishable food through the Panama Canal on a regularly scheduled basis. The ship's manifests, which detail everything being carried, must be presented to Canal authorities, just as the ship approaches the Canal. If the manifest is incomplete, the ship is sent to the end of the line, which means a 5-day delay. If this were to happen even once, it would be catastrophic for the company, so it was easy to justify a faster system that could deliver the manifest as the ship left the last port-of-call before the Canal.

Helping Clients Understand

The conventional wisdom is that only technology buyers apply this methodology. However, some sellers of products and services are beginning to adopt such an approach to help buyers compare a variety of complex offerings. Infonet Corporation (El Segundo, CA -- www.infonet.com) is among the leading vendors already using this methodology to help customers understand how to purchase international Frame Relay services. This is the vanguard -- a high-touch approach to sales and support of what are essentially commodity products with confusing options.

The greatest value from the methodology is that it gets managers of the business, applications, network and the vendor communicating towards a common set of objectives. While the studies can get very involved, in the end, customers wind up

buying services with a much clearer understanding of how various features impact the business. In fact, customers learn that sometimes they must change the application design or how it is used in order to achieve goals that cannot be delivered by any communications upgrade. The vendor may not sell that added feature this time, but the vendor will surely be asked to help solve the next problem. Helping a customer make an informed decision generates lasting value for the relationship.

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.