Managing Internet Performance

Net Forecasts – Peter J. Sevcik
BCR Volume 30, Number 3
March 2000

The conventional wisdom is that the Internet is unmanageable, because end-to-end performance is stitched together from thousands of ISPs, each of whom delivers only a small part of the vast Net. Since there is no central administration of the 'Net, there can't be any consistent performance.

But, many net-based applications will need guaranteed performance in order to succeed and, fortunately, the conventional wisdom isn't so much wrong as it is overstated: Application providers can take concrete actions to understand, improve and control performance to users.

Measure, Measure, Measure

The actions begin with an old truth: You can't manage what you don't measure. Begin by measuring service performance from the *user's* perspective. This can be done via specific tests using client-based agents such as RAPS from Quest Software (acquired Fog Light -- www.foglight.com) or Vital Suite from Lucent NetCare (which was acquired via Lucent's purchase of INS/VitalSigns -- www.ins.com).

Once a measurement system is initiated, however, the job is far from over; having a professional service monitor your site on an ongoing basis is a more comprehensive approach. Such services are available from Keynote Systems (www.keynote.com), Exodus (acquired Service Metrics -- www.servicemetrics.com), Mercury Interactive (www.merc-int.com), Optimal Networks (www.optimal.com), Visual Networks (acquired Inverse Network Technology -- www.inversenet.com) and others. These services can not only test from multiple locations, they enable you to avoid having to train a staff or operate a measurement system.

In general, these services follow the lead set by the Frame Relay measurement companies that are now used for insuring Frame Relay SLAs. However, the metrics measured are entirely different – e.g., page payload, single packet round trip time and total page load time -- and to date, *no* major NSP has decided to use one of these services to guarantee an Internet SLA.

Moreover, if you're not operating a public site, a general measurement service won't see your servers. But many of the service bureaus listed above will either sell you software so you can operate your own measurement system or they will set up and operate a private service for you.

Test and Fix

The second stage in managing Internet performance is proactive intervention. Uncovering the real reason for slow performance requires detective work -- sophisticated diagnostic tools and trained operators. Some of this work can be performed by the measurement service, but generally they only re-analyze data they've already gathered. Fresh data is likely to come only from your team.

The first diagnostic procedure is always a ping test to the two sites in question; simple reachability results and round-trip times are basic data. However, NSPs typically either discard or spoof ping tests and/or ICMP packets, because they don't want to handle ICMP requests nor do they want you to know the true performance of their networks.

Therefore, a more reliable test is an instrumented session to the application on the hosts; a lot can be learned from differential timing of packets within a legitimate session that the NSPs must support. You also can correlate internal network traffic data, such as peering circuit utilization and neighboring network data, with the incidents of poor performance. Finally, any performance data your partners share with you can be extremely valuable.

More often than not, your problems wont' be found within your primary service provider's network but, instead, with another backbone provider or a regional network. However, depending on your size, you can still have leverage. Imagine the reaction from the following phone call: "A major hosting client in Michigan, who represents 1 Tera-Byte per month of traffic into your network, is having difficulty with all traffic on your network in and out of New York. The other backbone provider we use gives us no difficulties into New York. Please improve your service to New York or we will have to terminate our relationship."

But to be in a position to make that call, you have to perform many tests yourself. Your NOC will have to know how to perform ping tests, generate synthetic traffic, gather data from multiple test points and perform tests across a coordinated group of test points. Many levels of sophistication will be required to accurately diagnose both the application-level symptoms and the network level cause. While the tools for such testing are often home grown, new companies are emerging that will soon provide instruments for more sophisticated diagnostics.

Move the Problem

It's also important to recognize that many applications simply will not meet their performance target if they are hosted in a single location. No matter how good the server farm or Internet connections, if users are located far away, it takes time to reach them.

Caching is the only way to reduce network delay in the Web page response time -- a typical Web page-load time can be reduced by as much as two-thirds across North America. The further the distance the greater the performance gain, and the time reduction can be even more significant if the page has a high level of content (graphics), complexity (turns) or both.

Caching has two primary benefits: Shortening network distance and geographic load balancing. In essence, caching moves the server closer to the user, which greatly simplifies the "test-and-fix" stage of Internet performance management.

There are two basic alternatives to caching. The first came in the form of proxy servers and browser-based caching (see "Designing a High-Performance Web Site", *BCR*, March 1996, p. 27). However, general-purpose caching server systems from companies like Infolibria (www.infolibria.com), Inktomi (www.inktomi.com) and Resonate (www.resonate.com) have eclipsed these techniques.

Better yet are transparent caching appliances from companies like CacheFlow (www.cacheflow.com), Cisco (www.cisco.com), Cobalt Networks (www.cobalt.com) and Network Appliance (www.networkappliance.com), which focus on content delivery without the costs and diversions a general-purpose operating system. Transparent caching generally delivers all of the content of a page at a cache hit rate of 50-60 percent.

The second basic approach is to forgo operating caching devices and rely on content delivery network (CDN) services, which are available from companies like Akamai (www.akamai.com), Cidera (formally Sky Cache, (www.cidera.com), Digital Island (acquired SandPiper – www.digisle.com) and iBeam (www.ibeam.com). Clients can push the content out to a wide range of geographically dispersed locations that would otherwise be cost prohibitive for any single enterprise. Most CDNs operate some form of reverse DNS look-up by which they can find the best server from which to deliver the content

Some content, however, isn't appropriate for CDN distribution. For example, the Akamai approach adds time to a page load as the browser participates in the closest cache discovery process. It then makes up the time by delivering large graphics from that nearby cache. However, if the page does not have large graphics, the total page time can be longer than getting the content directly from the source. Since the content provider has control over what will be in the caches, the overall result is fast delivery of some content at a hit rate of 90-95 percent. However, I expect things to get better for low graphic but highly interactive applications as a new breed of transaction delivery networks become available by the end of the year.

Enhance the Internet

Finally, the most direct approach to managing Internet performance is to put a device on one or both ends of the connection that will insure that the proper traffic gets assured service.

The most obvious approach is to back-haul a dedicated, high-speed circuit from the server to the client location. This is how most ASPs currently insure good performance to their customers. Bypassing the Internet may be practical for a few high-end, high-volume customer locations, but it defeats the basic goal of operating a general service on the Internet.

A more practical approach would be to ship a box – e.g., class-based queuing and traffic-shaping

devices from companies like Packeteer (www.packeteer.com) and Lucent (purchased Xedia, www.xedia.com) -- to key clients. These devices identify and prioritize traffic going across the access line, and some operate rate-shaping techniques that better match the TCP window size to the actual available bandwidth.

New software vendors like WWWhoosh (www.wwwhoosh.com) perform more sophisticated improvements by operating on both ends of the connection. Most of these new techniques require that plug-ins be installed to the browser or that special software modify the TCP/IP software at the client. The most promising techniques are TCP selective ACK (SACK) and HTML multithreading or page consolidation.

Integrated Solutions

While each of the above techniques provides performance management value, the optimal solution will involve operating several of the techniques in an integrated fashion. An early integrated product is QoSWorks from Sitara Networks (www.sitaranetworks.com) a performance appliance that provides real-time monitoring, class-based queuing, TCP rate shaping, packet-size optimization and transparent caching all in one box. Comparative tests from the Tolly Group (www.tolly.com) show significant performance enhancement and policy accuracy relative to some of the more narrow solutions from Packeteer and Xedia. Properly coordinated multiple techniques do add up to more than the sum of the previous parts.

All of the current techniques generally manage content against bandwidth within a static setting, but this is a long way from the Internet really requires. For example, most caches are sold to save on bandwidth rather than to improve user response time, but in the long run, new techniques and improvements will open the opportunity to provide dynamic transaction control of the user experience. Expect to see the ability to provide end-to-end SLAs between any two points across the total Internet

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.