Performance Issues Facing the World-Wide Web

Net Forecasts – Peter J. Sevcik
BCR Volume 29, Number 9
September 1999

While the Internet has changed considerably in recent years, the Web remains slow for conducting serious transactional business. Anyone who uses the Web probably has reached a similar conclusion, and it's clear that the WWW and 'Net won't fulfill their promise unless the right design decisions are made. What follows is a look at performance metrics for the Web and an analysis of why the Web performs the way it does.

I've been analyzing Web traffic since mid-1995, and I published my first set of results in *BCR* the following year ("Designing a High Performance Web Site," March 1996, pp 27-31). What I've learned since then is that basic Web-page delay is largely governed by two key parameters: Payload and Turns.

Payload refers to the number of bytes a server sends to a browser between clicking on a URL and seeing "document done" on the screen. *Turns* are the number of times a browser must exchange non-payload packets with any server on the 'Net; one query-response exchange constitutes a Turn, and TCP ACKs are not counted. A Turn represents the small pause while the browser waits for network round-trip delay plus server processing. It may not sound like much but the Turns add up.

Timing measurements fall into two categories: The best case and everything else. The best performance a site can provide to a user at a fixed distance over the 'Net is the *performance floor.* The floor performance changes very slowly over time as the page design changes, servers improve, more bandwidth is provisioned to the Internet, etc. When a portion of the 'Net gets congested or if distance increases, performance drops. The floor is the best performance delivered to the best-connected users. (To derive the performance floor index I use Application Expert from Optimal Networks (www.optimal.com), which supplies Payload and Turn statistics and predicts the floor performance.)

In 1995, a "typical" business Web home-page was 50,000 bytes of Payload and required 20 Turns. A "typical" business-to-business session -- dedicated access lines of 1.5 Mbps at the server and 256 Kbps at the browser along with an Internet path of 240 msec round-trip -- could deliver that home-page at a floor of 12 seconds.

To assess where things stand today, I solicited help from Keynote Systems (www.keynote.com), which I consider to be the true arbiter of performance – it downloads pages from the top 40 business Web sites four times per hour each business day from 50 measurement locations around the U.S. Keynote's has performance data dating back to December 1996, and its latest on-line testing tool -- Lifeline – lets you see the performance and page profile of any site for free.

Good News/Bad News

I profiled each of the Keynote Business 40 URLs -- *aka* the KB40 -- with Application Expert in order to compare the results of the two data gathering methodologies. To get a page performance floor, I calculate a payload-turn index (PTI); PTI is the total bytes sent by all the servers to the browser divided by 12,500 plus all the Turns divided by 4 (PTI=P/12,500+T/4).

Figure 1 shows the results of this analysis. The blue points are the floor values of the KB40 sites over the past ten months (Oct '98 to Jul '99). The red points are the medians of all the weeks' data that fell outside the floor range.

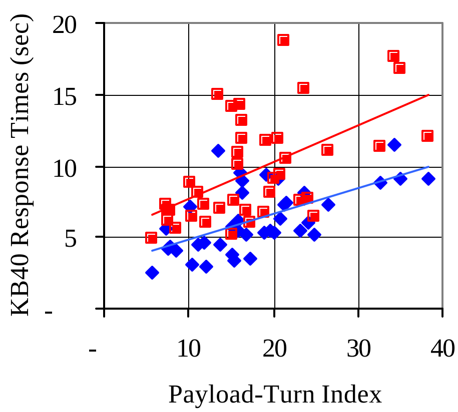

Figure 1 Using PTI to Predict Performance

Compiling my early data from 1995 with the Keynote data and the PTI statistics leads to a few conclusions, a mixture of good and not-so-good news. The basic business Web page has grown from 50,000 to 90,000 bytes, while the complexity of the pages as indicated by Turns has grown from 20 to 44.

Given those numbers, you'd think things would be much slower, but the floor has actually improved from 12 seconds to 6 seconds. There are a number of reasons why:

- In 1995, browsers did not multiplex or multi-thread TCP sessions, a capability that affects the Turns part of the PTI formula - the 1995 typical page PTI was 24.
- Older servers often use a small MTU and window in TCP, which adds overhead and hampers the ability to keep the connection fully utilized. The new server OS platforms and the browsers' multi-thread capabilities give Web designers much more leeway.
- Businesses are moving to 1 Mbps or higher access lines, while 256-kbps facilities were common in 1995.
- Both desktops and servers have more speed to process the highly intensive graphics on today's Web pages.

All of these factors combine to permit a page of about twice the size to load in *half* the time. The 'Net has been doubling overall Web delivery ability every 24 months.

However, and here comes the bad news, that may not be sustainable over the next 24 months, and here's why: Overall delay within the 'Net is significantly higher – round-trip delay has increased from 240 msec to about 370 msec.

Figure 2 shows where the total time is spent. Note that in 1995, computers, access line and the 'Net all contributed about equally to the 12-second floor. By 1999, however, the 'Net has become, by far, the largest component of the new 6-second floor.

Reducing 'Net Latency

Router hops are at the core of the delay problem, and the number of router hops to an Internet destination constantly surprises me. From my home office where I have a cable modem connection (MediaOne), there are at least 10 hops to all of the KB40 sites, and some are more than 40 hops away! There are four hops just to get out of my local MediaOne network, and that doesn't include all the layer-2 switches, ATM switches and firewalls that are in the path but do not report as routers.

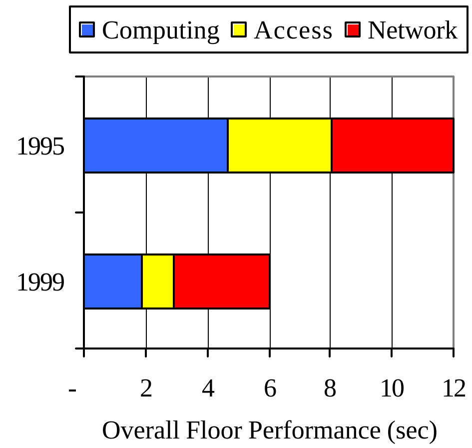

Figure 2 Where Does Net Delay Come From?

At each hop, it takes a 1,500 Byte packet 1.2 msec of serialization delay – at 10 Mbps! -- to get it into the router. Adding a mere 0.8 msec to process the packet means that hop costs 2 msec. At 40 hops, the packets are 160 msec away in round-trip before we start to add in that pesky speed-of-light delay for the convoluted path the poor packets must travel. (You can track the router hop component of delay by viewing the Internet Traffic Report published by the Andover News Network -- www.andovernews.com.)

The number of router hops is directly related to the complexity of access networks, backbone nodes and Web site networks. A typical major NSP node has feeder routers connected to core routers, which connected to ATM switches, which connected to more core routers, and all within the same building. Then you are off to other nodes, peering-points, networks and, finally, the destination site network, each with many routers. This is not how to run a fast network.

It's time for the service providers to start thinking about how to streamline their network topologies and the Internet in general, and the task won't be trivial. It will require that new, innovative technology be created, which can reduce store-and-forward hops within a physical node location. The growing complexity of Web pages and new highly interactive media will require getting closer to, not farther from, the speed of light.

How to Beat the System

There is no substitute for getting the content closer to the user, and today that equates to caching or content distribution. For example, it turns out that the few KB40 sites that have faster performance than the PTI predicted, use Akamai's (www.akamai.com) distributed content services. Akamai's service actually added Turns, but because the content was always delivered from a server closest to the Keynote measuring site, overall performance was much better.

Moreover, the Keynote servers are hosted inside major NSP nodes and connected by 10 Mbps to take the measurements. While it doesn't get any better than 10 Mbps, it often gets much worse. For example:

- The Keynote numbers are biased to the low side because they don't put failed hits into the performance averages.
- In the real world, users are not located within the ISP's POP, and they're rarely connected by a dedicated *broadband* connection.
- These numbers do not represent the experience of dial-up users and users outside the continental U.S.

Since things do get a lot worse, Web owners can take steps to ensure that each connection delivers every last ounce of performance. For example, a protocol called SNP from Sitara Networks (www.sitaranetworks.com) intercedes between the source and destination TCP sockets to fundamentally improve TCP performance. The Sitara protocol is targeted at Web sites that really must insure consistently high performance even when the Net is trashing or loosing packets.

A Faster Web Matters

So, is this all much a do about nothing? If the Web is merely a repository of stuff that people browse, why does it matter if performance for the well-connected improves from 12 seconds to six? There are two important reasons why Web performance matters.

First, the overwhelming majority of users do not experience 6 seconds. For example, my average for the KB40 ranged from 5 seconds at 2 AM to 19 seconds at 2 PM on a weekday. Moreover, our industry should not take any pride from the fact that it's taken four years for the performance floor to drop by 50 percent. Most IT elements improve much faster, and the Web has got to speed up its rate of improvement.

Second, a response time of less than 3 seconds is essential to maintaining full user concentration and productivity, so unless the Web speeds up, it won't be able to accommodate the many serious business applications that require enterprise-style response times. At the rate the Web is improving, the whole idea of Internet-hosted applications is in serious jeopardy.

A Web that consistently delivers pages in less than 3 seconds to the majority of its users is achievable, but not without major investments in new delivery technologies (fewer hops) and new Web content designs (fewer Turns). We know what to do, the question is does the industry have the will to get it done?

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.