Improving User Experience with Route Control

By Peter Sevcik and John Bartlett
July 2002

There is a common misconception that the performance of an application on the Internet simply cannot be controlled. Once an application is moved to the 'Net there is no value to measuring its performance because there is nothing that can be done about it. This flawed, fatalistic view usually results in one of two outcomes – a useful application is not placed on the 'Net or it is placed on the 'Net with disappointing results. But proactive management *can* control Internet performance leading to greater user satisfaction.

In this report, we study the range of performance that real users encounter on the Internet along with their perception of performance. We show that users experience a surprisingly wide range of response times, with some seeing times far worse than the average. We then show the potential benefit of using new route control technology to better manage performance and improve it dramatically for the most poorly served users.

Internet Performance

We all know intuitively that Internet performance can be fickle. The same Web page that responded quickly yesterday, or an hour ago, is suddenly very slow. We often don't know if the problem is on the Internet, at the Web server, or local to the user's PC. To answer questions about the range of Internet performance, some real data is required.

Internet Performance Data

As part of an ongoing assessment of Internet performance, more than 20 business-oriented Web sites instrumented their servers to determine the network delay for each user accessing their site. The data used in our analysis, collected over five days during August 2001, ran over eight different backbone carrier networks. Performance was measured by timing the TCP open sequence from the server's point of view. The time from the initial request for a TCP connection (SYN), until the time the client acknowledged the open connection (ACK) is a round trip time from the server to the client and back again. These times were recorded for all eight ISP links, not just the one selected by default Internet routing. Packet loss was also derived from the data.

The key benefit to this approach is that it measured performance across the complete Internet path of real users when and how they chose to access the sites. The raw measurement data was made available to NetForecast for analysis and assessment. The majority of the users accessing these business sites were connected to the Internet by broadband access (1 Mbps or faster). The detailed network data was also grouped into major classes of users by geography: US West Coast, US East Coast, and global.

Web Performance

The data regarding the performance of the Internet derived above is interesting but needs to be translated into Web response time as seen by a user. NetForecast has studied the fundamental interactions of the Web and how network performance affects Web page load time. This is well documented in "Understanding Web Performance" [1] describing the Keynote Business 40 Index (KB40) as a good basket of business-oriented Web sites. Applying the NetForecast Web performance model [1] to the network performance data shows a probability distribution of Web page load times as shown in Figure 1.

NetForecast Report 5062

©2002 NetForecast, Inc.

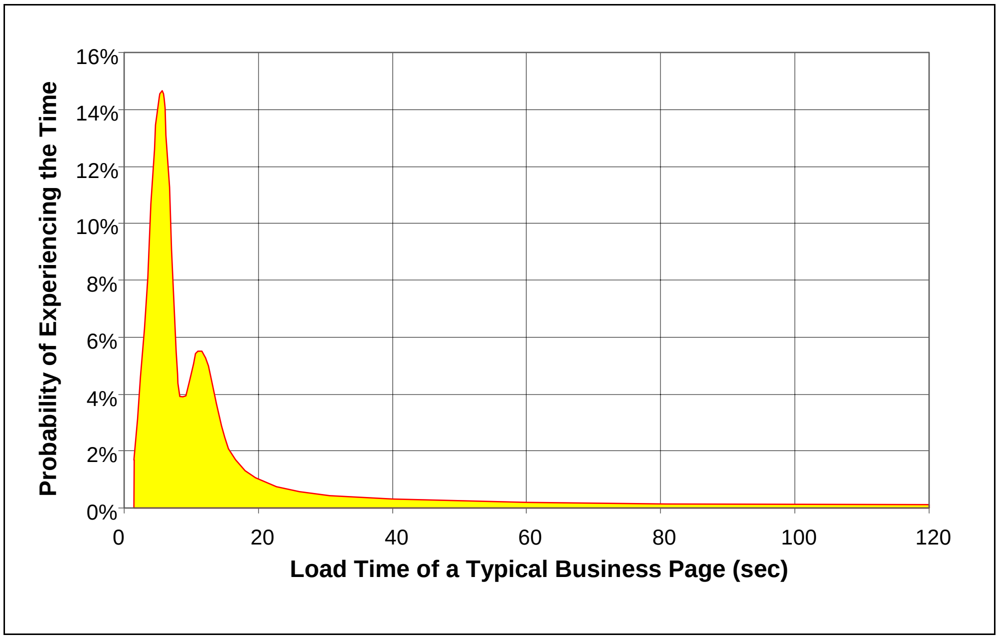

Figure 1 – Business Web Page Performance in North America

There are three significant observations that can be made by studying Figure 1. First, one can clearly see two peaks of user page load times at 6 and then 12 seconds. These peaks correspond to the users on the West Coast near the instrumented business sites (San Francisco bay area) and the next large population of users on the East Coast of the US. In fact, each of these significant user populations has its own curve that is aggregated in Figure 1. We can also conclude that there were about 3-times the number of West Coast users relative to the East Coast users.

Second, the median response time across the US was 6.9 seconds while the average was 12.0 seconds. Porivo also tested the KB40 in August 2001 across a more uniform population of broadband users throughout the US [1]. Applying the Porivo measurement methodology to the data in Figure 1 leads to an average response time of 10.9 seconds, which is a little slower than the 8.2 seconds measured by Porivo and reported in [1].

Therefore, the results of the two completely different approaches to measuring Web page response time correlate remarkably well. It must be pointed out that in August 2001, when both of these independent tests were being made to reconstruct the results of the KB40, Keynote Systems was reporting that the KB40 had a page load time of 2.5 seconds. See [1] for more details.

Third, and most important, there is a significantly long tail to the right on the curve in Figure 1. There are many users that see very long page load times. Figure 1 is cut off at 120 seconds on the horizontal axis in order to show the shape of the curve. In fact, there are points that reach more than 200 seconds!

Web Performance by Geography

Figure 1 and its analysis clearly show that there is no single value for a Web page load time. Averages and medians are important but they do not show the complete picture. To get a better understanding of the distribution of times any particular user can expect, we deconstructed the data into three separate groups, representing three geographies (local, far coast, overseas). We sorted through the data for each prefix, looking for the minimum value. A prefix represents one of the approximately 100,000 address ranges, or neighborhoods, on the Internet as defined by the Border Gateway Protocol (BGP) for Internet routing. The minimum value is predetermined by geographic distance, so each prefix was allocated to a geographic class based on this minimum round trip time (RTT) value. We further grouped the users into five equal population groups by ascending response time.

Figure 2 shows the bands of performance as seen by each fifth, or 20 percent of users in each major user population by distance from the server. The time scale is a log scale in order to better show the faster times. The time that represents each boundary is noted in seconds. As expected, the range of page load times for the middle 20 percent (yellow region) increases with distance. Also, the regional and national midpoints are on both sides of the overall US median of 6.9 seconds.

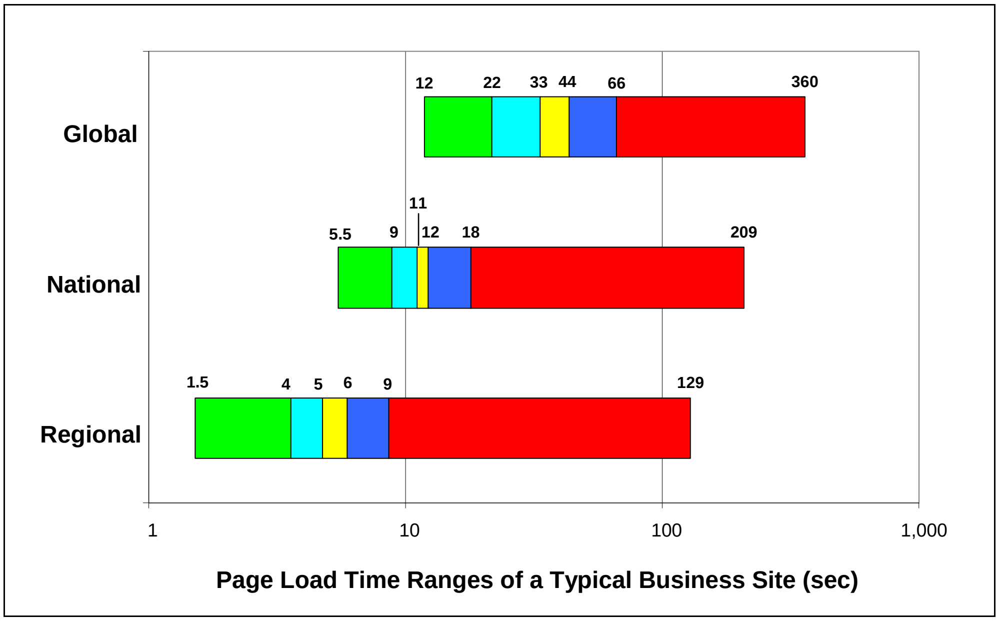

Figure 2 – Web Page Load Time Ranges in Different Geographies

There are several observations that can be made from Figure 2. Each fifth moving away from the center is progressively longer. The final segment representing the slowest 80th to 100th percentiles of users is very long. This is another view of the long tail in the distribution of time. For example, when regional users in the first 20 percent see a page load in 1.5 to 4 seconds, users in the final 20 percent must wait from 9 to 129 seconds.

The long tail is even more significant for the global users. One fifth of the global users had to wait for 1 to 6 minutes before the page completed! It must be pointed out that due to the methodology of the data gathered, we do not know how many of the users waited for the full 6 minutes. We do know that network conditions at certain times during the testing meant that some users would wait six minutes for a page to load completely.

Defining the User Experience

Much has been said about the effects of poor performance on the user experience. It is clear that users are affected by the performance of a Web page to varying degrees, but they do not think about which fifth of the user population they fall into.

There have been some interesting tests and measurements of the users' perception of application response time. The results of the research to date are summarized in "Understanding How Users View Application Performance" [2] where the case is made for three zones of performance that really matter to a user. The following applies the findings in [2] to the Web experience for the *typical* user in North America accessing a *typical* Web page in North America.

<u>Zone of Satisfaction</u> In the first zone of performance from 0 to 10 seconds, the user is not aware of the time it is taking to load the page [2]. The user is more aware of factors other than performance when asked about the page. These factors deal with the task at hand, ease of use, and other satisfaction criteria regarding the Web site. As long as a site can load in under 10 seconds, speed simply is not an issue in user satisfaction.

This single value is chosen because users are generally not aware of the physical location of a site in the Internet. Most users in the United States simply know that they are accessing a site in the United States.

<u>Zone of Tolerance</u> After 10 seconds, the user is aware of the fact that the page is loading over time. Time is now a factor in user satisfaction with the Web site. It starts as simple awareness and grows to frustration. If the page loads in a significantly long period of time, then the user is frustrated. Again Studies described in [2] show that the average US user would call page load time frustrating at about 40 seconds.

<u>Zone of Frustration</u> After 40 seconds, the user is past just waiting for the content to appear and becomes frustrated with the process. There is a strong likelihood that the user will stop the page load. It is also likely that the user will stop interacting with the Web site. The user may return at another time hoping that conditions will have improved. However, some users will likely not return to this site.

There are times when users will persevere with very slow sites. These are users who for some reason have placed a very high value on the content they will receive or have made a large investment in the outcome of the process. For example, if after much Web research, the user has found a report they are desperately seeking, then they are willing to wait to finally see the report.

Global Users

Users in most of the world outside North America are keenly aware when they are accessing information in the United States that it is "far away." The zone thresholds may be greater in these cases. The global users in this study are very far away – on the other side of the world from the servers. However, in order to maintain a conservative view regarding the tolerance of users, the thresholds of 10 and 40 seconds are applied to the global user population as well.

Route Control

Route control is a new network function that optimizes how traffic traverses the Internet. Route controllers are used at an enterprise or hosting site where more than one connection to the Internet is provided, either via different carriers or by one carrier to more than one major POP location. Figure 3 shows a typical multi-homed Web content provider, with connections to backbone ISPs A and B.

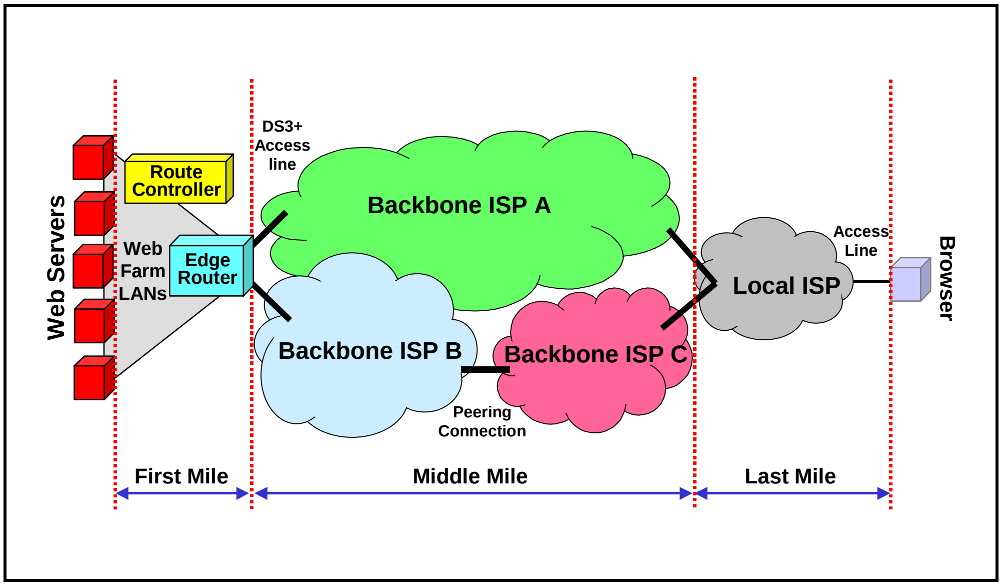

Figure 3 – Alternate Routes for the Multi-Homed Web Site

The Border Gateway Protocol (BGP) is the default mechanism that decides how traffic will flow through the Internet. BGP chooses the path based on the fewest number of Autonomous System (AS) hops. An AS represents a domain of control, usually a single organization (carrier or ISP). Routers in each AS converse with each other through a BGP session and forward routing information. This information insures that the link between the two ASs is operational, and indicates the network distance from each AS to each destination prefix on the Internet.

If a link is broken, BGP updates will occur, and traffic will be forwarded via a different route. However BGP does not pay any attention to latency or packet loss on these connections. If the connection through carrier A is running slowly, but is still alive and has the fewest number of AS hops, then carrier A is the choice.

The route controller may be implemented as an appliance that lives at the enterprise network edge, or it may be a service provided remotely as described in “Route Control: Optimize Multi-homed Connections for Performance, Load and Cost” [3]. In either case, the route controller monitors the performance of traffic from active portions of the Internet to the enterprise site, noting round trip delay, packet loss, and the carrier that was used, along with the remote IP address or prefix.

As the route controller builds up its database, it is able to see that traffic to certain Internet locations has better performance through carrier A than through carrier B, whereas other remote locations may do better through Carrier B than through Carrier A. The route controller updates the forwarding table of the enterprise edge router on a prefix-by-prefix basis to favor the carrier providing better performance for that prefix. The update rate varies by vendor and is defined by policy. Of course the route controller can only affect the portion of the network where alternate paths are available.

Impact of Route Control on Performance

NetForecast wanted to understand the value of improved performance by the best carrier choice relative to the current default BGP choice in terms of the number of users affected, and the amount of performance change that would occur. To do this, we modeled a nearly perfect route controller in the following way.

Best Case Performance

The measurement data set consists of about 400,000 hits across eight carriers. For each client prefix and carrier combination in chronological time, we maintained a running average of the TCP open time. This running average represented the current performance on that carrier for that prefix. Then we identified the expected performance for each hit using the BGP carrier choice, and the expected performance using the best carrier, that is, the carrier with the lowest running average time.

Our simulation represents a best-case route controller, in that it makes all its routing updates immediately. If a hit comes in from prefix X on carrier A with excellent performance, we immediately start routing traffic destined to that prefix over carrier A. Real route controllers have to limit the number of routing table updates to the edge router because of the performance impact on the router.

Our simulation is limited by the data set available. For some prefixes we have many hits occurring, which will make our moving averages quite accurate, while others had few hits. The results for these prefixes were either optimistic (if a fast hit came in first), or pessimistic (if a slow hit came in first). However this is likely to be the case in the real world as well, and perhaps the optimism and pessimism balanced out in our combined results. Integrating this information will then establish a ‘best path’ before the user traffic arrives. The impact to performance when using the best path is shown in Figure 4.

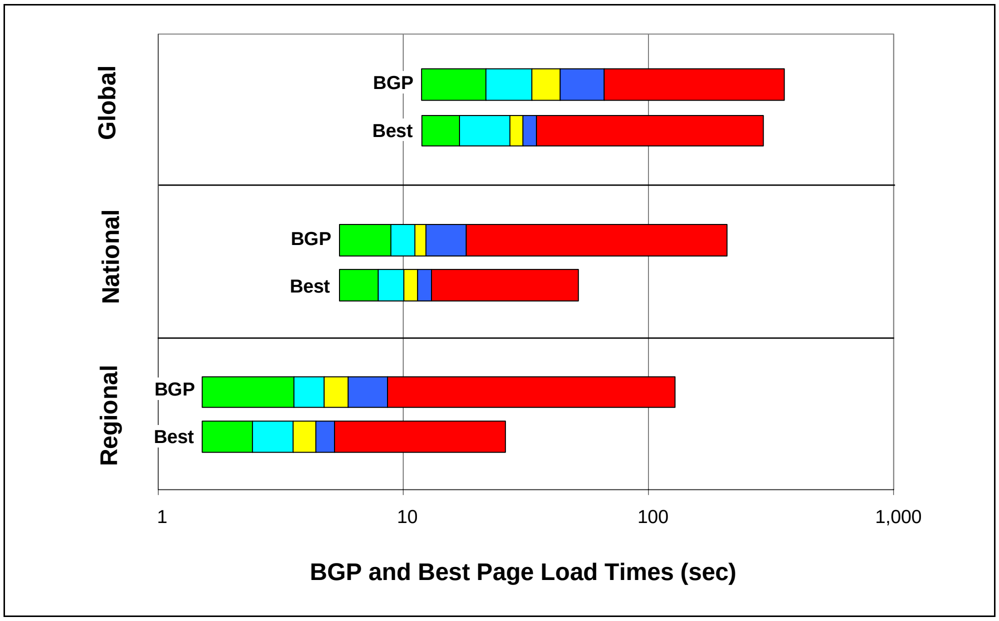

Figure 4 – Comparing BGP and Best Performance

The results in Figure 4 are grouped in the same methodology as in Figure 2 where each color represents the range of performance for each fifth of users. Figure 4 shows that each fifth of users shifts to a faster time from BGP to Best. Note that performance for the most poorly served segment of users, denoted by the red bars, shifts significantly to the left. The horizontal scale is a log scale where small changes as seen by the eye are actually significant. For example, the 80^{th} percentile time shift by geography is: regional, 8 to 5 seconds, national, 18 to 13 seconds, and global, 66 to 35 seconds.

Real route controllers have an advantage over our approach by being able to implement active measurement in addition to passive measurement. Our performance measurements come from passively watching Web page requests initiated by Internet users. If an enterprise knows where important users reside, their prefixes can be actively tested with measurement probes sent through each available carrier.

Applying the User Criteria of Performance

When the user performance criteria described above is applied, the results of the two approaches becomes more meaningful since this is how the user will perceive the performance. Table 1 shows the number of users who experience performance within the ranges believed to be critical for interaction with a Web site.

As would be expected, many (83 percent) users see satisfactory performance when the server is in their region. However, the majority (70 percent) of users merely tolerate performance when the server is on the opposite coast of the United States. If users are equally distributed across the US, then about half (53 percent) are satisfied, leaving half that need improved performance.

Table 1 – Number of Users in Each Performance Category

		Satisfied	Tolerating	Frustrated
Regional	BGP	83%	14%	3%
Same coast USA	Best	94%	6%	0.0%
National	BGP	22%	70%	8%
Across USA	Best	31%	68%	0.2%
Global	BGP	0%	31%	69%
Other side of globe	Best	0%	51%	49%

Nearly Eliminating Frustrating Performance

The results of using route control to change route selection from the default BGP choice to an optimized 'best' carrier approach are dramatic. Most regional and national users that had seen frustrating performance are no longer subjected to such excessive page load times. The number of satisfied users improves by 10 percent in all US geographies. Even on a global basis, the number of users seeing frustrating performance is reduced by 20 percent.

If a Web server is located on one of the US coasts and has a user population uniformly distributed across North America, then 17 percent – one in six – of the users will experience a better performance zone.

However, due to the fundamental limits of propagation delay, in order to truly solve the problem for global users, the Web service must be mirrored to a location closer to the users. Mirroring to at least the continent needing service is a minimal step. Route control can then be used to improve the user experience on each continent as described above for the US users.

The Route Control Challenge

The question for route controller vendors is how close to this 'best case' scenario we have outlined can they get? Vendors use a broad mix of approaches to determine the performance of each path. Some use active measurement while others use passive traffic monitoring or a mix of both. How often are the routers updated with new forwarding information? How accurately can the route controller move the prefix for a group of users to the best path? Is the update prioritized to optimize the most important destinations first?

Enterprises must evaluate the approach of route controller vendors before making a purchase decision. They must determine whether an approach meets the needs of the enterprise in providing performance consistency, load balancing and cost reduction. Moreover, they must determine how these three factors interact, since achieving all three all the time may not be possible. Finally, enterprises must decide whether there are sufficient controls to permit setting desired policies and adequate reporting to document whether those policies are delivering the desired results.

Summary

It is prudent policy to purchase Internet access from more than one ISP as protection against potential service disruptions. Route control can leverage this protection into a performance-enhancing asset.

Clearly there is a huge potential for improving performance for the most poorly served users. The distribution of times that users experience has a very long tail, and a significant number of the users are the outliers in that tail. Mangers of Web sites and other applications on the Internet should view this approach the same way manufacturers view defects, with the goal of reducing them to as near zero as possible.

For the first time, enterprises can find and improve performance seen by outliers rather than ignoring them. Reclaiming just half of the frustrated users as revenue generating customers for a Web site with revenue of $10 million or more annually can easily justify the purchase of route control technology. ¤

References

1. "Understanding Web Performance," by Peter Sevcik and John Bartlett, October 2001, NetForecast Report 5055.
2. "Understanding How Users View Application Performance," by Peter Sevcik, July 2002, Business Communications Review.
3. "Route Control – Optimize Multi-homed Connections for Performance, Load and Cost," by John Bartlett, January 2002, NetForecast Report 5057.

Peter Sevcik is President of NetForecast and a leading authority on network technology. He specializes in performance analysis and holds the patent on application response time prediction. Peter has contributed to the design of more than 100 networks including the Internet. He can be reached at peter@netforecast.com.

John Bartlett is a Principal at NetForecast, and has 24 years of experience designing at the chip, board, system and network levels, with a focus on performance. John led the team that built the first VLAN implementation, one of the first ATM switches, and he is a leading authority on the behavior of real-time traffic on the Internet. He can be reached at john@netforecast.com.

NetForecast develops customized analytic models to determine the business value of new technologies.

www.netforecast.com