NetForecast

Akamai Helps Improve Web-Enabled SAP Performance

By Peter Sevcik and Rebecca Wetzel
May 2008

Remember the old FedEx package delivery slogan, "When it absolutely, positively has to get there overnight"? Well—delivering business-critical applications on time is analogous to delivering business-critical packages on time, and when it comes to delivering SAP over the Internet, many factors can impede on-time delivery.

As an application delivery network service provider, Akamai is able to overcome many Internet performance impediments to help keep SAP users productive regardless of their distance from the SAP server. In this report we describe Internet-specific impediments to SAP performance and how Akamai addresses them, and we document how Akamai's technology can improve the experience of Internet-connected SAP users around the world.

Akamai's application delivery network dramatically shrinks the virtual distance between the SAP user and the server. Based on NetForecast's real-world modeling, in general an Akamai-enabled user up to 9,000 miles from a server experiences performance similar to that of non Akamai-enabled users under 3,000 miles from a server, and an Akamai-enabled user between 9,000 and 12,000 from the server experiences performance similar to that of non Akamai-enabled users between 3,000 and 6,000 miles from the server.

Business Importance of SAP Performance

Unlike most software applications, SAP is business critical by nature. While other applications support a gamut of uses ranging from 'nice to have' to 'mission critical', SAP applications are always used to run the business—the lifeblood of the company.

SAP applications are as ubiquitous as they are business critical. With over 43 thousand customers in more than 25 industries in 120 countries, SAP is the world's largest business software company, and the third largest independent software provider.

NetForecast's research has determined that SAP users' tolerance for long response times is low, especially for highly repetitive tasks. If SAP performance suffers, so does the business, therefore, performance expectations and requirements for SAP are more stringent than for other business applications.

As the lifeblood of the company, there is no such thing as SAP "lite". SAP is always in the business spotlight, so it "absolutely, positively has to" perform well, even for users who are far from SAP servers.

SAP-over-Internet Performance Threats

SAP performs efficiently over the Internet compared to most other applications, but it cannot escape a number of common Internet performance threats. Chief among these are distance and congestion. All of these factors conspire to slow SAP performance over the Internet and cause it to be unpredictable.

NetForecast Report 5093

Understanding Application Response Time

To understand how congestion and distance impair application performance, it is important to understand key factors influencing application response time (R). These include: application turns (round trips), application payload, network bandwidth, network round-trip time (RTT), server compute time (Cs), and client compute time (Cc).

The following formula shows how these factors combine to determine a user's response time for a given task (see the *Glossary of Terms* at the end of this report for detailed formula definitions).

$$R \approx \frac{\text{Payload}}{\text{Bandwidth}} + \text{AppTurns}\left(\text{RTT}\right) + \text{Cs} + \text{Cc}$$

Distance as a Threat

Distance hurts performance because the further a user is from the application server, the longer it takes packets to traverse the distance. This is sometimes referred to as the "speed of light problem". The time added due to distance (a.k.a. latency) can add up, especially if an application requires many round trips to complete a task.

SAP is generally centralized rather than distributed, which means that remote users often must traverse long distances before reaching the server. Long user-to-server distances bump up round-trip times and slow application response times for users.

Congestion as a Threat

Network congestion can stem from a variety of sources to consume available bandwidth and slow user response times. Common congestion sources are high bandwidth applications such as streaming audio and video, or "flash crowds" caused by surges of users simultaneously vying for the same limited bandwidth resources.

SAP is only one among the throng of applications contending for Internet bandwidth, and that bandwidth can be in short supply—especially at Internet service provider (ISP) peering points. SAP traffic must compete on an equal footing with all other applications in an Internet-resource free-for-all.

Without protection, SAP can readily be starved of essential bandwidth, and as the above response time formula shows, when SAP bandwidth is squeezed, user response times rise.

Akamai Approach

Akamai's application delivery network service—Akamai Web Application Accelerator—speeds and enables predictable performance for dynamic, interactive Web applications by applying an array of performance optimizing technologies, including:

- Dynamic Mapping
- SureRoute Route Optimization
- Akamai Protocol

The Akamai Protocol addresses the inefficiencies of HTTP and TCP to dramatically reduce the number of round trips (AppTurns) over the Internet.

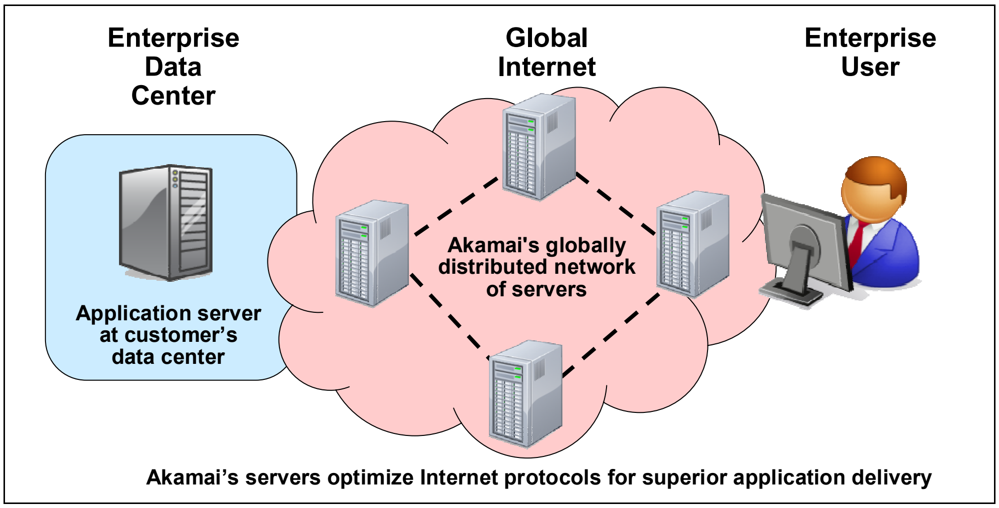

Figure 1 – Akamai Service Architecture

Application requests and responses between users and the origin server are sent over Akamai's application delivery network shown in Figure 1. Dynamic mapping directs a user's request to the optimal Akamai edge server, where route optimization finds the fastest, most reliable path to the server to retrieve the application content. Application retrieval and delivery is then optimized using a combination of TCP optimization, dynamic compression, route optimization and application turn reduction as the acceleration "levers" shown in figure 2 below (see Glossary of Terms for definitions).

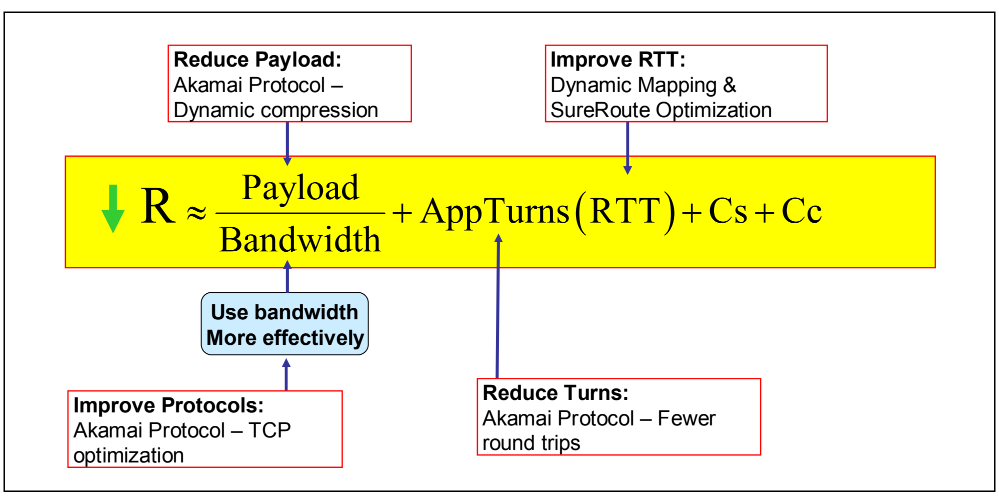

Figure 2 – Acceleration "Levers" Akamai Uses to Improve SAP Performance

NetForecast Performance Model

NetForecast has studied application performance for over a decade and has embodied that knowledge in a detailed model accounting for the performance factors that affect user response times. The NetForecast model uses a multi-variant Monte Carlo analysis technique that applies real-world distributions to parameters in the response time formula.

Among the factors accounted for in the model is the inescapable fact that round-trip time is directly proportional to distance divided by the speed of light—therefore, the longer the distance, the longer the round-trip time. The model also takes into account the fact that packet loss escalates with distance because the risk of congestion increases with higher network hop counts. The model can also show effects of bandwidth contention among applications, as well as the consequences of other aspects of network stress such as high network latency and loss.

NetForecast has found that it is unhelpful to compare raw response time values for each application because task definitions and judgments about acceptable task completion times vary widely for each application type. For example, what might be a painfully slow task time for one application task type might be entirely acceptable or even excellent for another.

Therefore, to equalize the view of performance across applications and network scenarios, we apply an application performance index methodology called Apdex, developed by the Apdex Alliance as an industry standard for measuring and assessing user satisfaction with application performance (see the *How Apdex Works* sidebar).

SAP application tasks are highly repetitive and designed for short query responses, and SAP users typically require a target time of two seconds to remain productive. So we assigned the SAP over the Internet application itself an Apdex target of two seconds (Apdex T=2). For contextual purposes, this contrasts to typical target response times of one second for SNA, nine seconds for Microsoft MAPI, and 12 seconds for Microsoft CIFS.

The model applies a distribution of server response times centered on a 50 millisecond average. NetForecast has seen SAP applications that achieve performance at the data center as modeled here. But we have also seen SAP implementations that must dig through extremely large databases for many of the task responses and take much longer to deliver the task reply. In such cases, the Apdex target T must be scaled up appropriately. The relative performance and conclusions drawn in this report would not change by applying a larger T value appropriate to longer server response times.

How Apdex Works

Apdex provides a numerical measure of user satisfaction with the performance of enterprise applications. Apdex reports are arrived at using a four-step method.

Step 1 – Select Target Time

Define a target response time of T seconds. Response times between 0 and T are considered satisfactory by the users within a business scenario.

Step 2 – Measure Performance

Measure performance and place each measurement into one of three user-perceived performance zones: satisfied (Sat), tolerating (Tol), or frustrated. The satisfied zone is defined in Step 1. Frustrated times are above F seconds, where F=4T. Tolerating times are between T and F seconds.

Step 3 – Apply the Apdex Formula

Sum the incidents of response times that fall into each zone, and divide by the total number of measurements using the formula (the sum of frustrated incidents is not in the numerator but is included in the denominator):

$$Apdex = \frac{Sat + Tol/2}{Total}$$

Step 4 – Show Results

Show the Apdex Index value (on a scale of 1 to 0) together with the corresponding threshold T. The Apdex Scale is subdivided into performance quality regions ranging from excellent to unacceptable.

In this report, we describe how T is selected for each application (Step 1). For brevity, the details of Steps 2 and 3 are not shown. Results of Step 4 illustrating the Apdex Quality Regions are shown in Figure 3.

Apdex is an open standard developed by the Apdex Alliance. Visit www.apdex.org to learn more.

SAP Performance Results for Akamai

NetForecast applied real data to a use case scenario to create a model for SAP performance over the Internet with and without Akamai's service, and the results showed dramatic improvement for global users with the Akamai Web Application Accelerator service turned on.

The Use Case Scenario

The NetForecast model for SAP performance was built around a multi-step business scenario in which a production user requires excellent performance. In the scenario, the user processes 50 sales orders, each of which requires four tasks with a target time of two seconds. In addition, each sales order requires downloading a sale summary file with a target time of nine seconds.

The Results

Figure 3 shows the effects of server-to-user distance on the performance of the above SAP use case over the Internet. Without assistance, SAP baseline performance dives into the poor range within 6,000 miles of the server, and into the unacceptable range within 9,000 miles of the server. With the Akamai Web Application Accelerator service turned on, however, the NetForecast model shows that Akamai stretches good performance an additional 7,000 miles from a baseline of 4,000 miles out to 11,000 miles from the server.

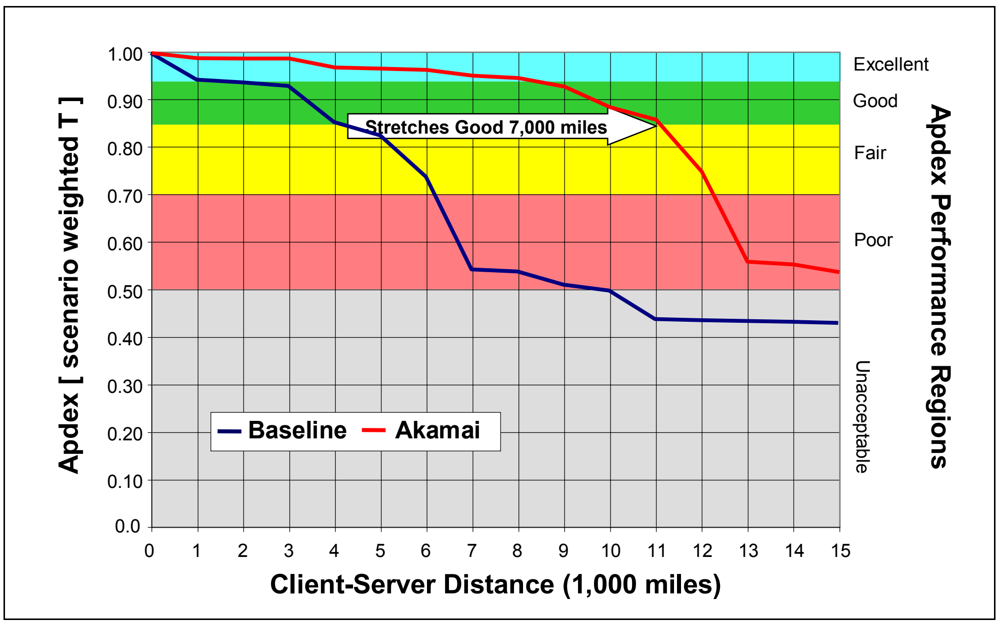

Figure 3 – Effects of Akamai Web Application Accelerator on SAP Performance

But what does it really mean to say Akamai "stretches" good performance? To show Akamai's effect on SAP performance experienced by geographically distributed users, we mapped Internet distance performance zones using the Apdex quality regions shown in Figure 3 and a user's physical location relative to the data center housing the server. Because Internet backbone circuits run on a predominately east-west axis, each zone is wider than it is tall. Figure 4 shows the baseline (no acceleration) performance zones for this SAP business use case for users accessing an SAP server in Chicago. The Chicago-based server delivers poor or unacceptable service to a very large population of potential global application users as Figure 4 shows.

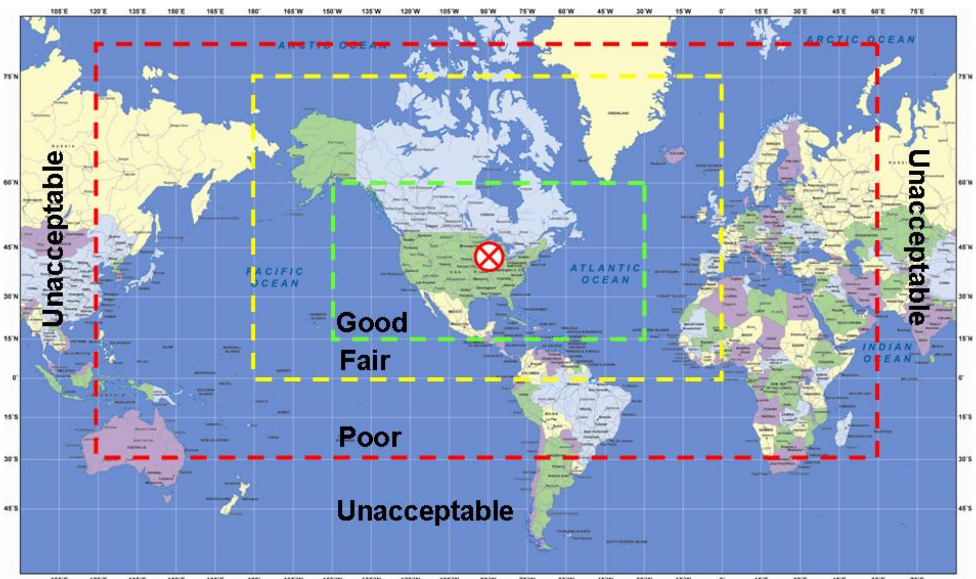

Figure 4 –Baseline Internet Distance Performance Zones (Server in Chicago)

Figure 5 shows how Akamai "stretches" the boundaries of good performance for Chicago-homed users to 11,000 miles east-west from the server, and stretches fair performance to 12,000 miles east-west—in essence covering the entire globe. In this scenario the Akamai Web Application Accelerator service eliminates the risk of poor or unacceptable performance.

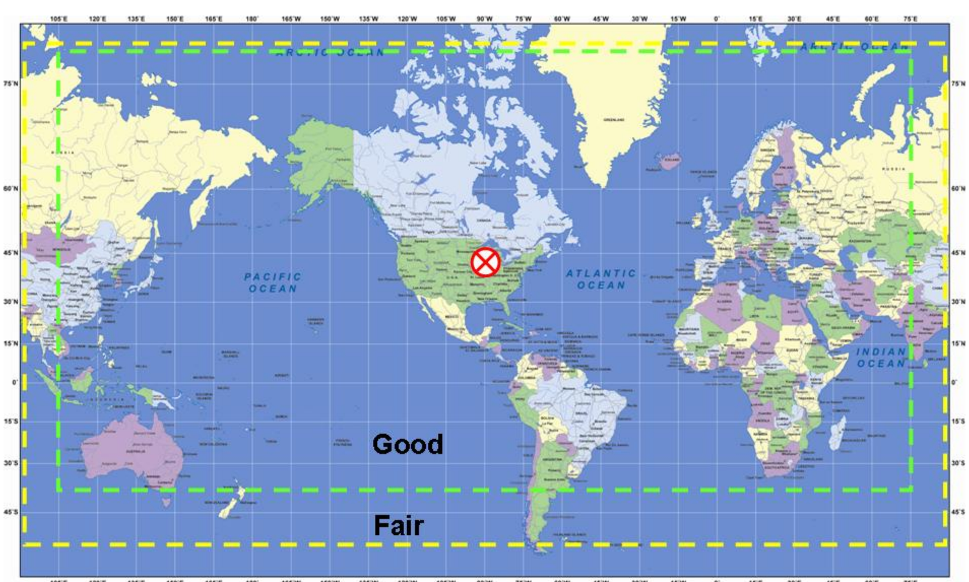

Figure 5 – Akamai Internet Distance Performance Zones (Server in Chicago)

The Akamai performance levels shown in Figure 5 generally hold for any data center location within the continental United States.

Without Akamai's assistance, all users in Asia and most users in Europe experience unacceptable performance for all US data center locations. The difference between supporting overseas SAP users without and with Akamai's Web Application Accelerator is significant enough to change a business strategy from failure to success.

Apdex:
Apdex is an open performance reporting standard defined by the Apdex Alliance. See: www.apdex.org

Conclusion

If you use SAP over the Internet and your users are outside your geographic region, you will need a solution like Akamai's to get acceptable performance. As the adage goes, your mileage may vary from what our results show. But we can say with confidence that your mileage will improve when you apply Akamai's Web Application Accelerator to keep your SAP users happy and business productivity high regardless of where they are located in the world.

Glossary of Terms

AppTurns are the application client-server software interactions (round trips) needed to generate a user-level system response or task. Turns are separate form and in addition to TCP two-way interactions (e.g., ACKs). The user is not aware of turns.
Application Turn Reduction reduces the application turns by intercepting many original client-server transmissions locally and re-transmitting all the content in a single block.
Bandwidth is the minimal bandwidth (bits per second) across all the network links between the user and the application server. The slowest link is typically the user's access line to the network. Useable bandwidth of the link may be reduced by the effects of conflicting traffic (congestion) and protocol efficiency (e.g., TCP window).
Cc (Compute Client) is the total processing time (seconds) required by the client device.
Cs (Compute Server) is the total processing time (seconds) required by the server(s).
Dynamic Compression is applied to data "on the fly" to reduce payload.
Frustrated Zone Any application response time above F in which the user is very negatively affected by response time.
Payload is information content (bytes) that must be delivered to/from the user's device.
R Response time is defined as the elapsed time (seconds) between a user action (e.g. mouse click, enter, return) and the system response (client, network, server), so the user can proceed with the process.
RTT is the round-trip-time (seconds) between the user and the application server.
Satisfied Zone The range of application response times between 0 and T in which the user is not affected by the response time.
T is the target threshold (seconds) which defines the boundary between satisfied and tolerating performance zones.
Task Each user-application interface interaction that requires a user entry and an application response.
TCP Optimization increases throughput by matching the transmission rate to a constrained bandwidth access line, and/or overriding TCP window limits on high bandwidth but long latency paths. Pre-emptive data receipt acknowledgements maintain high throughput to speed data from the source, and ramp up TCP transmission rate more quickly by bypassing TCP's 'slow start' function.
Tolerating Zone The range of application response times between Satisfied and Frustrated in which the user is negatively affected by response time.

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

Rebecca Wetzel is an Associate of NetForecast and a veteran of the data networking industry with unparalleled inside knowledge of the Internet service and product markets. She works with network product vendors and service providers to develop and implement product strategies. She can be reached at rebecca@netforecast.com.

NetForecast helps enterprises understand and improve the performance of networked applications.

Additional information is available at: www.netforecast.com