Monitoring Next Generation Web Applications with AlertSite DéjàClick

By Peter Sevcik
June 2007

Smart business and technology choices can have a negative unintended consequence: Visibility into the end-user experience is lost. Many changes in the way applications work over a network must be accompanied with commensurate changes in how application performance is monitored. This report describes these issues and presents a new tool from AlertSite called DéjàClick that directly addresses the problems of monitoring next-generation applications.

Next Generation Web-based Applications

The web is undergoing a transition from a collection of autonomous websites to an integrated computing platform serving web-based applications to end users. This phenomenon, often referred to as "Web 2.0," is not really a change in any web standards but rather a shift in how standards are applied. The real change is that the web application is being deconstructed into separate elements which can reside on the primary website or be distributed among many locations on a network. Each element can then be supplied by a business partner or simply a friend whose website is added to your mash-up. Within an enterprise this change is driven by SOA and virtualization.

These changes in how applications are being built are supported by many well known benefits such as leveraging new standards, the ability to build applications more quickly, and making the user interface a rich, intimate coupling to the application. However, the new style of application design introduces some unintended consequences that negatively impact performance management.

The role of the browser is shifting from a passive client to a dynamic client, as shown in Figure 1. In the passive client mode, all of the presentation logic is loaded to the browser in the base-page portion of every web page. The browser then dutifully gets each page element and paints the pages as instructed. In the dynamic client mode, the presentation logic has moved to the browser where software operating above the browser makes choices based on what the user desires.

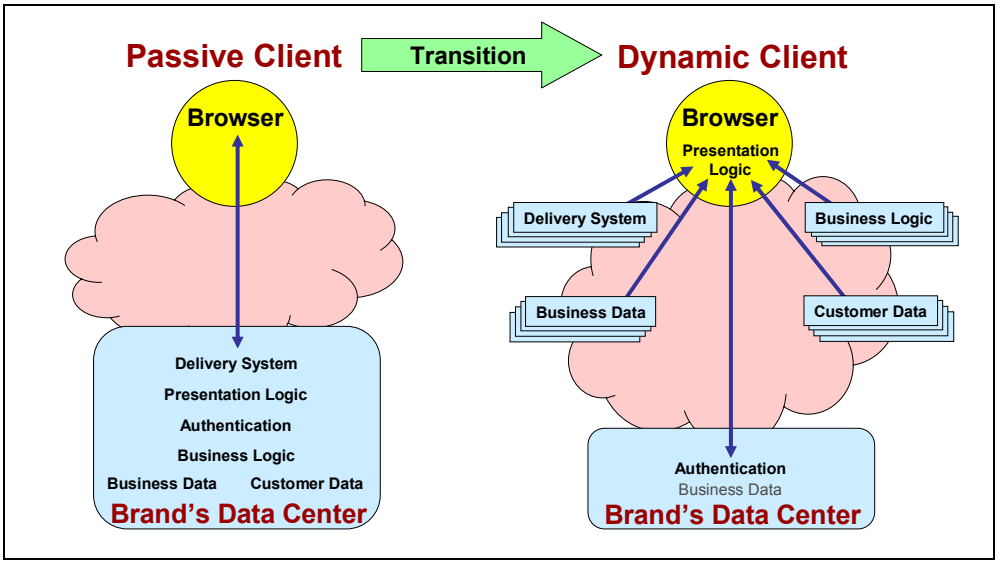

Figure 1 – Shift in How the Web Works

NetForecast Report 5087

The dynamic client is a level of software that operates above the browser and within the browser window. Common examples of a dynamic client are Java scripts, Ajax, and Flash. In each case the user is primarily interacting with the dynamic client that is operating within a browser tab. In effect, the dynamic client is yet another level up the protocol stack above the browser's layer 7, which would make it a new layer 8.

A very important difference in the dynamic client model is that the data required to fill out elements of the web page no longer come from a single location. Clearly some of the contents may come from an alternate delivery system like a CDN. Recently, however, business logic, customer data, and even the business data that is needed to complete a transaction are being supplied from third party sources.

The browser performs the final assembly of the web pages based upon a combination of rules defined by the primary web server, dynamic client, and the user. (The user does not directly set rules, but rather just moves a mouse or makes choices on the screen.) Furthermore, there are many aspects of the application that operate on the desktop that do not interact with any server on the network. The dynamic client is autonomous software that is the new true user interface, rather than the browser which is relegated to a supporting role.

The user, with a richer set of controls over the browser and dynamic client, has greater impact on how the web page will look. The user can set policy over cookies, ad-blocking, and the ability to run Java or ActiveX. Just within a browser, the user can have multiple tabs open, each of which may be updating data or images in the background. All of these settings are unique to a user's desktop and may have an impact on just how a web-based application works or even what it displays.

Losing Web Visibility

The reality of Figure 1 is that the traditional places where the user experience is monitored no longer apply.

Most enterprises monitor web performance by instrumenting the data center. After all, that is where all of the components of the web page originated. Furthermore, all the components that were generated in application servers and database servers were traditionally funneled through the web server to the browser. So instrumenting the web server from web log analysis to sophisticated business logic analysis provided a lot of information. But as parts of web pages are delivered from locations outside the data center, this approach sees fewer elements of the full page.

Another good place to monitor the user experience is in the data path between the web server and the browser, accomplished by adding a passive data gathering appliance. There are several vendors of super-fast packet sniffers coupled with sophisticated software that reconstruct all of the HTTP and HTML into each user interaction. These real-user measurement tools provide a real-time view into the user experience. But the single blue arrow in Figure 1 has morphed into many arrows, causing the data center appliance to see less and less of the content.

Finally, there are synthetic measurement approaches that run a user script of the application from agents on the Internet. This has worked well as a way to check on the completeness and timing of a web page or sequence of pages. The synthetic agent's accuracy is a function of how well it can replay *all* of the actions of a real user. Unfortunately, most measurement services are not upgrading their agents to the level and complexity of a dynamic client. AlertSite is addressing this need, which will be described below.

Dynamic Client Examples

Dynamic Clients are appearing in many subtle ways. Often, users are operating an application or interacting with a website unaware of the new paradigm. Here are some everyday examples.

Electronics Retailers

The typical on-line store has many more items available than the brick and mortar store. This is certainly true of electronics retailers. For example, a major electronics retailer can easily have 1,000 televisions in its catalog. A buyer can search and compare these televisions by a large number of technical details such as screen size, display type, overall product dimensions, weight, inputs, outputs, and remote features. This is just a sampling of the more than 50 technical details available. And they change as each of the suppliers drops or adds a model.

The retailer cannot keep the 50,000 details on all the televisions they sell current. So the retailer lets the details about each product come directly from the manufacturer. Critical business information pertinent to a sale appears on the user's browser without the retailer involved.

A major on-line store can easily have 100 such relationships where the content presented to the user comes directly from the manufacturer.

Newspapers

Many newspapers offer an on-line "front page" that is customizable by the reader. In fact, the on-line newspaper is an integration of many sources. Again, key elements of the reader's web experience such as the weather, stock ticker, horoscope, advertisers, social networks, and even the newspaper's content search window come from third parties.

It is often simpler and safer to have such content come directly from the source rather than trying to store, update, and coordinate the content from the third party. A typical newspaper can have such direct data feeds originate from dozens of outside sources.

Banks

Most banks provide an on-line account management capability. The information regarding a customer's bank balance, checks cashed, and on-line payments comes from the bank's computers. But many banks provide additional services as they try to be the single place to go for a wide range of financial needs.

Some banks provide an interesting bill presentation service. For example, customers can link their electric bill directly to their bank's on-line interface. This is not the ability to have the electric company automatically withdraw a payment from the customer's checking account each month. Instead, this offers customers the ability to review their electric bill from within the confines of an on-line banking session.

When the customer seeks information from the electric company, that portion of the bank's web page is handed over to the electric company. The session, along with the bank's passing of the user's authentication credentials, is handed to the electric company. Without additional steps, the customer is now in a direct dialogue with the electric company systems. The customer can review the electric bill and choose to pay some, all, or none of it. The customer's payment choice is relayed to the bank and the customer may continue with the bill payment process.

The critical point is that for a period of time, the customer's secure SSL session was passed to a third party. The bank had no view of what transpired between the customer and the electric company. Similar secure hand-offs occur for purchasing insurance or switching to the bank's brokerage partner.

Your Brand is on Top

In the above examples, the primary brand was the retailer, newspaper, and bank. When some portion of the web page was being supplied by one of the partners, the brand on the top of the web page was that of the retailer, newspaper or bank. If something goes wrong such as content is not delivered, or performance slows to a crawl inside the partner's window, the user still thinks he is talking to the primary brand. The real danger is that your brand's reputation and revenue is at a heightened risk because a portion of the session is no longer under your control. Higher risk with less visibility is a recipe for disaster.

How AlertSite DéjàClick Works

DéjàClick is a new product from AlertSite in Coconut Creek, Florida. DéjàClick is a free software tool coupled with the TrueUser Service which is a fee-based monitoring system. Let's look at the two offerings separately.

DéjàClick

DéjàClick is a free tool bar addition to the Firefox browser (see download link at the end of this report). It is a sophisticated desktop-based agent technology that operates on your desktop. There is a critical difference between DéjàClick and other desktop agents. Instead of running as an agent on top of the operating system simulating network events, it operates on top of the browser or the dynamic client simulating a user as shown in Figure 2. This is critical to seeing how the application operates within the confines of a real browser.

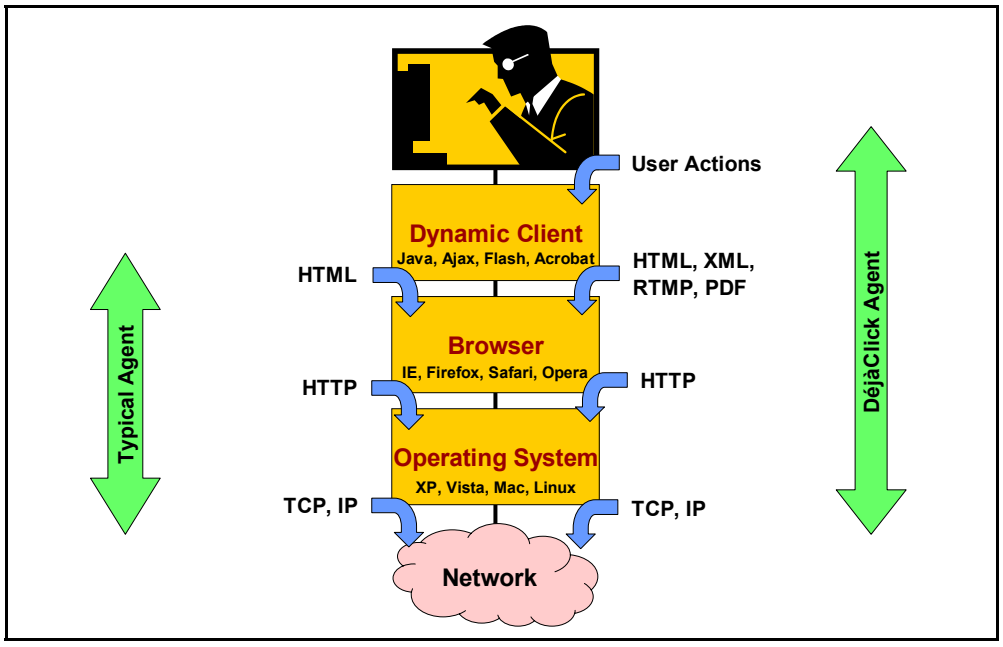

Figure 2 – Where Agent Software Operates

Programming the agent is easy since it includes a transaction recorder that follows the user's click stream to automatically build scripts. You can record a simple interaction or a complex session. The resulting script is a detailed outline of each user action, client event, or browser event. The user actions include such "internal to the desktop" events as "mouse hovering over button" that would not be seen by recording below the browser.

During initial playback, algorithms within DéjàClick notice if there are variations between how the script was recorded relative to how it is playing. It includes a pattern-

matching capability to correlate the content delivered with the content originally recorded. The application tester can modify the script to account for a wide range of conditions such as user think-time, network activity, cookie caching/hiding, form-filling requirements, etc. The built-in Replay Advisor guides the tester through these choices and recommends changes to the script replay parameters.

Once a script meets the needs of the tester, it can be saved locally, turned into an autoplay Super Bookmark, or uploaded to AlertSite for repetitive testing from their network.

TrueUser Business Transaction Monitoring Service

DéjàClick scripts can be loaded onto the network of agent servers operated by AlertSite. The service is being rolled out to more than 30 locations where AlertSite has servers. The DéjàClick script operates just as it did when it ran on the tester's desktop. It is simulating user actions that trigger dynamic clients and Firefox browsers operating worldwide.

The scripts operating on the remote agents have an adaptive playback capability in order to continue to properly work in the presence of small changes in the application, content, or software environment. For example, it automatically detects and adjusts to cookie-based form-filling. It also automatically adjusts to asynchronous activity and changes in the browser environment (e.g., an action causes a new tab to open).

The remote agents do not add any latency to the software interactions, browser functions, or network transaction. This is a true measure of real user performance as seen from the AlertSite monitoring locations.

There are several controls that can be set on the remote agents to affect the adaptive playback nature of the scripts. For example, it can be set for differences between first- time and subsequent interactions by a user. The agent management interface uses a self-service model to set up and change a sophisticated remote monitoring system. However, AlertSite provides a support staff that can help set up a monitoring system within the TrueUser service.

The scripts can operate at any interval up to a rate of every 5 minutes. The data gathered resides in a database which can be viewed under a variety of reports. The data can also be exported to perform off-line analysis.

DéjàClick is a work in progress. AlertSite is adding capabilities as new dynamic clients and features emerge. Today, it supports a wide range of conditions that are found in media-rich applications.

How to Use DéjàClick

Global monitoring with simulated users can detect network anomalies, server problems, and delays in receiving content from partners. Traditional synthetic agents, passive appliances, and server log technology can't discover these issues if the website uses complex rich-media design. Enterprises should use DéjàClick and the TrueUser service to perform the following management functions.

New Application Quality Assurance

No Internet application is static. It is important to test any changes before they are rolled out to the operational environment. You can use DéjàClick to test changes to the site, supporting software, or content partners. This is a valuable tool for ensuring that the user experience is not hampered by the change.

Failure Discovery

DéjàClick provides a foundation for setting thresholds on various technical aspects of the experience. The TrueUser system has many ways to set alarms that alert administrators if

basic links stop working, content matching fails, etc. Alarms can be used to discover subtle changes in how a part of the user experience has changed

Problem Diagnoses

Diagnostics determine which part of the complex system of users, networks, and servers is creating a performance problem. With so many "moving parts" to the problem, diagnosis is really a matter of deductive reasoning. DéjàClick helps eliminate non-offending parts of the system in order to narrow the search for the troublemaker.

Correlation Analysis

TrueUser measurement data can be used over time to develop a performance baseline. Analysis of the data will discover if performance is consistent across geographic regions. The data can also be correlated across multiple service points such as various Internet service providers, hosting services, or content delivery networks. In fact, using DéjàClick is one of the few ways to independently verify the service being provided by a CDN to the parts of the world that matter to your business.

Recommendation

Business success is tied to properly managing core business operations. While focusing on core operations, however, it is likely that you will leverage third parties on the Internet to play supporting roles. It is also likely that in order to remain competitive, you will add sophisticated media-rich interfaces to your on-line user interface.

Your monitoring system must follow these trends. You cannot assume that the third parties are supporting your goals without verifying performance. You cannot assume that just because it works at headquarters, it is working where the real users are located.

New, complex, media-rich or dynamic client applications must be monitored with the next level of testing that is provided by DéjàClick. Since DéjàClick is available at no charge, you should encourage your IT staff to download the tool and run simple tests of your application. Building sophisticated scripts is easy and free. You can fire them off as needed from the Super Bookmarks. Once you are familiar with DéjàClick you should subscribe to the TrueUser network service to implement processes that continuously track the user experience.

Dynamic client applications often sneak into a website with little fanfare. It is likely that some of your web measurement data is already missing the full picture. Don't become complacent because page load times have improved when, in fact, that result is based on measuring only a portion of the user's experience. Things may not be as good as they seem. You should at least check to make sure.

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.

AlertSite is a leading provider of web performance measurement, systems monitoring and security vulnerability scanning products that ensure a customer's critical web-based services are always available and running at peak performance. Founded in 1998, AlertSite is based in Coconut Creek, Florida and has over 2,000 customers worldwide. More information is available at:

Web: www.alertsite.com
DéjàClick free download: www.dejaclick.com
Phone: 877-302-5378

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

Additional information is available at: www.netforecast.com