

Monitoring the Web 2.0 User Experience with AlertSite DéjàClick

By Peter Sevcik and Rebecca Wetzel
January 2010

You can't improve what you don't measure—and what you measure matters. In the case of dynamic Web 2.0 applications, what you measure must reflect the *real* user experience—which can only be captured from within the browser. Why is that? Because dynamic Web 2.0 applications often obscure visibility into performance, so they must be accompanied by changes in how application performance is monitored. This report describes these issues and profiles a tool from AlertSite called DéjàClick that addresses the problems.

Passive Is Passé

The Web is transitioning 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 web applications are being deconstructed into separate elements which can reside on the primary website or be distributed among many network locations. Each element can then be supplied by a business partner whose website is added to your mashup. Within an enterprise this change is driven by service-oriented architecture (SOA) and virtualization.

These application design changes are being spurred by new standards, and the desire to build applications more quickly as well as to couple the user interface more richly and intimately to the application. However, Web 2.0 application designs adversely affect performance management.

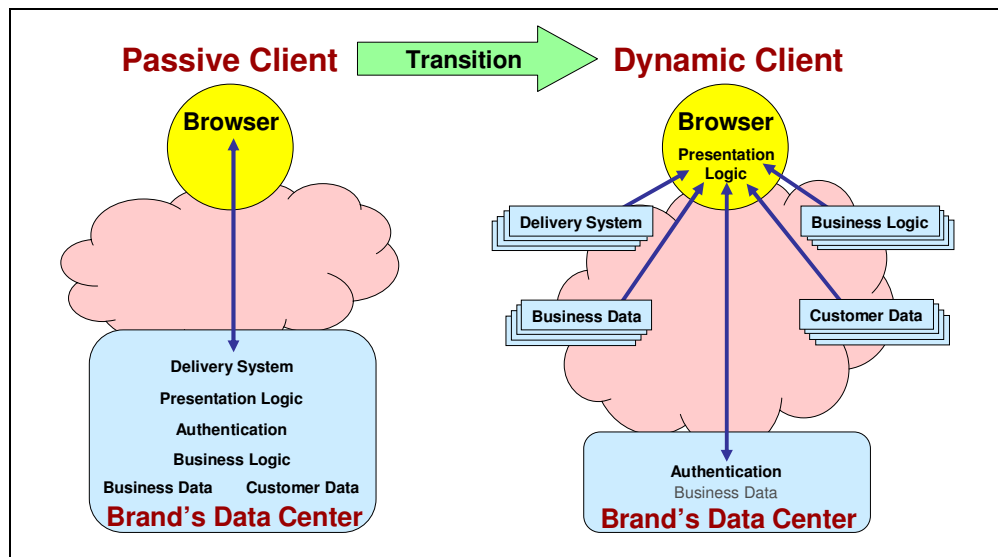


Figure 1 – Shift in How the Web Works

The browser is shifting from a passive to a dynamic client, as Figure 1 shows. In passive client mode, all 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 dynamic client mode, the presentation logic moves to the browser where software operating above the browser makes choices based on user actions.

The dynamic client 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 operating within a browser tab. In effect, the dynamic client is one rung up the protocol stack from the passive browser's layer 7, which would make it a new layer 8.

A key 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 source. Some content may come from an alternate delivery system like a content delivery network (CDN), but more recently, business logic, customer data, and even the business data needed to complete a transaction are supplied from third party sources.

The browser assembles 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 simply moves a mouse or makes choices on the screen.) Furthermore, many aspects of the application operate on the desktop and do not interact with any server on the network. The dynamic client is autonomous software that is the new true user interface, relegating the browser to a supporting role.

With a richer set of controls over the browser and dynamic client, the user can now customize how the web page looks. 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 affect how a web-based application works or even what it displays.

Losing Web Visibility

The reality of Figure 1 is that monitoring the user experience from the traditional locations no longer works.

Most enterprises monitor web performance from the data center because before Web 2.0, that was 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 more web page elements are delivered from outside the data center, this approach sees fewer elements of the full page.

Another place to monitor the user experience is in the data path between the web server and the browser, by adding a passive data-gathering appliance. Some vendors couple super-fast packet sniffers with sophisticated software that reconstructs all of the HTTP and HTML in each user interaction. These measurement tools provide a real-time user experience view, but the single blue arrow in Figure 1 has morphed into many arrows, causing a 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 works well to check the completeness and timing of a single web page or sequence of pages, but a 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 do not have agents that are sophisticated enough to accurately monitor a dynamic client. AlertSite addresses this need, as we explain below.

Dynamic Client Examples

Dynamic Clients are appearing in many subtle ways. Most users operating a web-based application or interacting with a website are unaware of the new paradigm. Here are some everyday examples.

Online Retailers

Online retail sites are providing increasingly rich and diverse applications that are hosted or “fed” by third parties. Here are just a few among a myriad of examples.

Manufacturer-supplied Product Information: The typical online 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 adds or drops a model.

No retailer can keep some 50,000 details on all the televisions they sell current. So the manufacturers directly feed details about each product to the retailer. Critical business information pertinent to a sale appears on the shopper’s browser with no retailer involvement.

A major online retailer can easily boast 100 such relationships in which user-presented content comes directly from the manufacturer.

Sophisticated Design Features: Increasingly sophisticated Java and Ajax applications and software-as-a-service (SaaS) offerings hosted by third parties are pushing the envelope of what is possible for online retailers. For example, it is now possible to try on clothes in a virtual changing room so prospective buyers can see how clothing will look before purchasing. With dynamic new interactive applications, businesses such as printers and sign makers now enable customers to design everything from business cards and wedding invitations to yard signs and company banners to their own specifications.

The highly interactive aspects of these features are implemented in software running on top of the browser. The web server provides the basic options available and receives the final selection made by the consumer. The performance of sophisticated design process that is critical to the user experience can only be monitored on the user’s PC.

Newspapers

Many newspapers offer an online “front page” that is customizable by the reader. In fact, the online newspaper is an integration of many sources. 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.

Financial Service Providers

Most financial institutions provide an online account management capability. Information regarding a customer’s account balance, online payments, and checks cashed comes from the financial institution’s computers. But many financial institutions provide additional services as they try to be a one stop shop for a wide range of financial needs.

Online Bill Presentment: Some financial institutions provide online bill presentation services. For example, customers can link their electric bill directly to their financial institution’s online interface. This is not the ability for the electric utility to automatically withdraw a payment from the customer’s checking account each month, rather, customers

can review their electric bill from within the confines of their financial institution's online session.

In some cases the server changes. For example, when the customer seeks information from the electric utility, that portion of the financial institution's web page is handed over to the electric utility. The session, along with the financial institution's passing of the user's authentication credentials, is handed to the utility. Without additional steps, the customer is now in a direct dialogue with electric utility 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 financial institution, 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 financial institution had no view of what transpired between the customer and the electric utility. Similar secure hand-offs occur for purchasing insurance or switching to the bank's brokerage partner.

Financial Tools: Many financial institution websites are adding a growing array of informational tools. These include such things as mortgage refinancing and retirement savings plan calculators. New tools are appearing constantly to help with such things as financial performance pattern matching, trading based on trigger parameters, and trading based on pre-set formulae. Such tools are often delivered to users by software-as-a-service (SaaS) applications in conjunction with companies that specialize in such software tools.

Financial Information: Investment firms often aggregate information from a variety of sources to aid customers in making well-informed trading decision. This information, which can include such things as key financial performance statistics and financial analyst ratings can originate from sources like Standard and Poor's, Morningstar, Thomson Reuters, etc.

Healthcare Providers

Healthcare sites are providing increasingly rich and diverse applications that are delivered under the healthcare provider's primary brand.

Health-related Tools: To encourage health awareness and wellness, healthcare websites are adding a growing array of informational tools. These include such things as calculators for health-related metrics ranging from body mass index and ideal weight, to target heart rate and pregnancy due dates. New tools are appearing constantly to help with such things as medication tracking and management, disease risk assessment and diagnosis, and chronic disease management. Such tools are generally delivered to users by software-as-a-service (SaaS) applications that are hosted by companies that specialize in such software tools.

Health Record Access: Increasingly, healthcare websites offer consumers online access to their health history and current status information. In many cases such back-end capabilities are hosted not by the website owner, but by a third party who feeds the requested data to the consumer.

Physician Directories and Hospital Comparisons: Many healthcare websites provide the ability to find doctors by location, specialty, and other criteria. This information is often provided by partners, and does not come from the healthcare website itself. The same is true of incorporating capabilities such as WebMD's Hospital Advisor so health insurance websites can enable customers to compare hospitals using a variety of criteria such as length of stay and complication rates.

Clinical Trial Recruiting: Many hospitals are involved in clinical trials, and provide the ability for the public to learn about the trials. This information is often available through partners.

Your Brand is on Top!

In the above examples the primary brand is yours. When some portion of the web page was being supplied by a partner, the brand on the top of the web page was that of the website owner. 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 free user-experience monitoring tool from AlertSite of Coconut Creek, Florida, USA. DéjàClick is coupled with AlertSite's TrueUser™ Service which is a fee-based monitoring system. Let's look at the two offerings separately.

DéjàClick

DéjàClick is a free Firefox browser toolbar add-on (see download link at the end of this report). Its sophisticated agent technology operates on your desktop. Unlike other desktop agents that wedge between the operating system and the browser to simulate network events, DéjàClick operates as a simulated user on top of the browser or the dynamic client as Figure 2 shows. This vantage point is essential to seeing how the application operates within the context of a real browser.

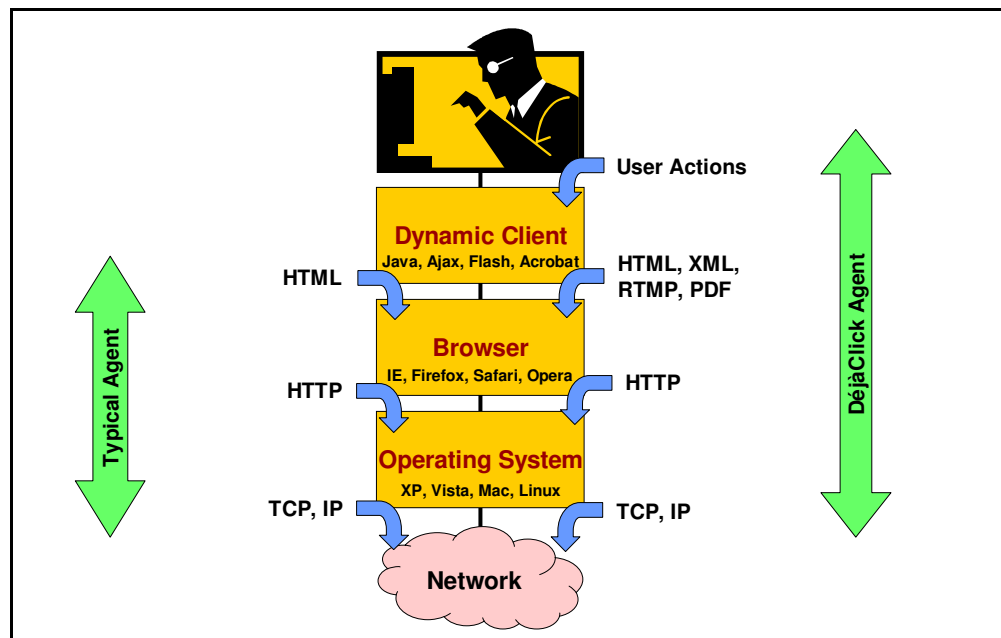


Figure 2 – Where Agent Software Operates

Programming the agent is easy because a transaction recorder automatically builds scripts by following the user's click stream. 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 never be seen when recording beneath the browser.

AlertSite recently enhanced DéjàClick with TrueScreen™ technology that monitors and records all user interactivity with Adobe Flash objects as well as all Ajax-based applications. TrueScreen monitors all Flash media user activity, including keystrokes, mouse clicks and drags, as well as “mouse overs”, to accurately record user-Flash object interactions.

During initial playback, algorithms within DéjàClick identify variations between how the script was recorded and how it is playing—and pattern-matching algorithms correlate the content delivered with the content originally recorded. The application tester can modify the script to account for a wide range of factors such as user think-time, network activity, cookie caching/hiding, form-filling requirements, etc. A 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 auto-play Super Bookmark, or uploaded to AlertSite for repetitive testing from their network.

TrueUser Business Transaction Monitoring Service

DéjàClick scripts can be loaded onto the TrueUser network of agent servers operated by AlertSite in more than 50 locations. The DéjàClick script operates just as it did on the script developer’s desktop, but it now simulates user actions that trigger dynamic clients on Firefox browsers operating worldwide.

The scripts operating on the remote agents have an adaptive playback capability that enables them to work properly in the event of small changes in the application, content, or software environment. For example, the adaptive playback capability automatically detects and adjusts to cookie-based form-filling, as well as 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 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.

You can set several controls on the remote agents to affect the adaptive playback nature of the scripts. For example, you can set them to simulate differences between first-time and subsequent user interactions. The agent management interface uses a self-service model to set up and change a sophisticated remote monitoring system, and AlertSite also provides support staff to help set up a monitoring system within the TrueUser service.

The recent addition of Java-based custom scripting permits an even more powerful way to add logic to the script. For example, the script can select information from a web page that is then entered into a form on the page or on subsequent pages. This is a way to dynamically generate tests that are context sensitive. For example, the script can discover what item is on sale and then proceed to make inquiries specifically about that item.

The scripts can operate at any interval up to one minute. The data gathered is stored in a database which can be viewed in the form of a variety of reports. The data can also be exported for off-line analysis.

DéjàClick is a work in progress. AlertSite has plans to add capabilities as new dynamic clients and features emerge.

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 Discovery

NetForecast research shows that just over half of all application performance problems in enterprises with best-in-class performance management practices are discovered by end users rather than by monitoring systems—while in enterprises with worst-in-class practices, users discover fully 85 percent of performance problems. The True User system can be a valuable aid to identify performance problems before users are affected.

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 and improve mean time to repair (MTTR).

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.

Load Testing

DéjàClick scripts can be written to randomize navigation events to reflect the actions of a real population such as to mimic the actions of first-time users or repeat users, and to reflect a range of user “think times”. By realistically load testing your site to reflect the actions of real users you can more accurately identify areas of stress vulnerability.

Recommendation

Measure what matters, which in the case of Web 2.0 applications is the real user experience that can only be captured from within the browser. Dynamic client applications often sneak into a website with little fanfare, and when they do your web performance measurement methodology is probably missing the full picture. It's easy to become complacent when page-load times look good, but if you only measure a piece of your user's experience, things may not be as rosy as they seem.

Complex, media-rich or dynamic client applications can only be monitored with comprehensive testing like that provided by DéjàClick. Since DéjàClick is free, NetForecast recommends that you download it and give it a test drive. Building sophisticated scripts is easy. You can fire them off as needed from Super Bookmarks. Once you are familiar with DéjàClick, we recommend you subscribe to AlertSite's monitoring services to continuously track the user experience.

Peter Sevcik is a principal of NetForecast and is a leading authority on Internet traffic, performance and technology. Peter has contributed to the design of more than 100 networks, holds the patent on application response-time prediction, and pioneered many performance management techniques including Apdex. At BBN Peter played a pivotal role in the development of the Internet. He can be reached at peter@netforecast.com.

Rebecca Wetzel is a principal 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.

AlertSite is a leading provider of Web performance management solutions that ensure a customer's Web-based services and applications are always available and running at peak performance. More than 2,300 customers worldwide rely on AlertSite to understand and enhance the end-user Web experience, including AT&T, Honda, Ryder, Symantec and Purina Mills.

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

NetForecast and the curve-on-grid logo are registered trademarks of NetForecast, Inc.