

Citrix® NetScaler® Web 2.0 Push: The End of Server Sprawl?

Abstract

On April 14, 2009, Citrix announced the release of a new NetScaler 9 enhancement, NetScaler Web 2.0 Push. This solution has the potential to dramatically reduce the massive server loads generated by some Web 2.0 technologies by enabling servers to service up to 10 times more users than is possible without this technology. Delivered as an enhancement to NetScaler 9, Web 2.0 Push will be available via upgrade to current NetScaler customers with valid maintenance agreements, and also be provided on new NetScaler 9 appliances.

This ENTERPRISE MANAGEMENT ASSOCIATES® (EMA™) Impact Brief details these enhancements. It also discusses the reasons why this is a milestone technology that could well have a similar impact on Web 2.0 delivery as the load balancer had on the delivery of traditional Web applications.

Web 2.0 in the Data Center and NetScaler Web 2.0 Push

Web 2.0 has literally changed the face of the Web, from static “pages” to colorful, interactive “sites” that are slick and responsive. At a 30,000 foot level, Web 2.0 is a loose description for a broad approach that implies a second generation of the Web and includes a new class of Web applications often called Rich Internet Applications (RIAs). Examples of the technologies supporting RIAs include AJAX, discussed in this paper, Adobe Flash (explore at: <http://www.adobe.com/flashon>) and Microsoft Silverlight (explore at: <http://memorabilia.hardrock.com/>). RIAs draw the user in as a participant and have become powerful tools for social networking, competitive differentiation, and online commerce.

From the end-user perspective, one hallmark of Web 2.0 applications is that although they are hosted on the server, they have the richness of functionality typically associated with desktop applications.

From the end-user perspective, one hallmark of Web 2.0 applications is that although they are hosted on the server, they have the richness of functionality typically associated with desktop applications. Web 2.0 is often also associated with user participation, personalization (such as FaceBook at: <http://www.Facebook.com>) and collaboration (such as Wikipedia at http://en.wikipedia.org/wiki/Main_Page). Many Web 2.0 sites have become so popular and so pervasive that their annual growth is measured in 100s of percentage points.

While Web 2.0 has been enthusiastically embraced by users and businesses alike, it has taken its toll in terms of IT provisioning and support costs. Traditional Web applications consumed considerably less server resource. Servers had to “work” only in response to a user request, and when a server was not serving up a request, it maintained no connection to the user PC. Web 2.0 totally changes the game, since AJAX and similar technologies require additional user connections and server processing.

The bottom line is that delivering this rich content generates an exponentially heavier server load than traditional Web applications.

The bottom line is that delivering this rich content generates an exponentially heavier server load than traditional Web applications. Although an entire paper could be written on the server-side implications of Web 2.0, a short list includes the following:

- *Rich Functions:* With functions like “zoom,” a common capability in applications like Google Maps, for example, each zoom generates a torrent of content delivered from server to browser, increasing server processor and memory load as well as bandwidth. Multiply a single user performing one zoom in Google Maps by hundreds of users performing multiple zooms and it is easy to see why server loads have increased as Web 2.0 has proliferated (for an idea of the impact on Google, take a tour of one of Google’s container-based data centers at: http://blogs.computerworld.com/video_take_a_look_inside_a_google_datacenter?source=NLT_PM)

- *“Persistent” or “Long Lived Connections” (the key problem Citrix solves with this announcement):* One of the biggest changes from traditional Web to Web 2.0 applications is the idea of persistent or long-lived connections. In traditional Web applications, a user request to access a Web page retrieves the page from the server and, once the page is delivered, the server-side work is done.

In contrast, many Web 2.0 applications maintain an ongoing connection between the server and the user PC to push data to the user until the user leaves the page and/or site. This adds significant server load because, in addition to the initial loading of the Web page, the server must also manage connections to each user, as well as ongoing interactions between server and browser. Adding users, particularly those that stay connected for long periods of time, means that the server has one more connection to manage, and the number of users a single server can handle is limited by its processing power.

- *AJAX and Reverse AJAX:* AJAX is a group of technologies that changes the game in terms of application responsiveness, and streaming. Reverse AJAX and Comet are implementations on the server to push data to the user. These are used together to enable the delivery of low-latency content and use persistent connections to do so. Client Polling and Server Pushing are examples of the ways in which these technologies are used, and are briefly explained below. Citrix NetScaler Web 2.0 Push addresses the “Server Pushing” aspect of Reverse AJAX, Comet, or Server Push applications.

- *Client Polling:* After the page is loaded into the browser, the client application continues to query the server for updates.

- *Server Pushing:* Via a persistent connection, the server sends new data to the browser as it becomes available. Yahoo’s (<http://www.yahoo.com/>) financial site is likely built using this or similar technology. As the user views the site, stock price changes are displayed in near real time, and font color changes from green to red according to whether the price is increasing or decreasing. Server Push adds to the interactivity and timeliness of the Web experience, but again at considerable cost to the server in terms of both processing and user connection management.

Despite the considerable expense of Web 2.0 technology, companies have made these investments because of their compelling business value. Users have come to expect the graphics-laden, interactive experience that Web 2.0 provides, and companies are racing to fulfill user expectations with a positive experience that can pay off in customer satisfaction, better brand equity, and increased revenue.

What Citrix is Announcing

Citrix's new NetScaler Web 2.0 Push technology solves one of the big problems that Web 2.0 creates for CIOs and IT Operations—the problem of delivering rich Web applications at a reasonable price. This enhancement offloads the management of persistent connections to Citrix NetScaler 9.

For readers unfamiliar with NetScaler, it is a multi-purpose appliance designed to optimize and protect Web traffic. NetScaler can act as a firewall and perform data compression to improve network performance and similar tasks. Its most familiar role is as a load balancer, ensuring that the least busy Web server is first in line to process a new Web request.

Citrix Web 2.0 Push offloads much of the “work” currently performed by servers in support of Web 2.0 applications to the NetScaler appliance.

Citrix Web 2.0 Push offloads much of the “work” currently performed by servers in support of Web 2.0 applications to the NetScaler appliance. Where a single server had to manage and maintain concurrent connections to each Web client, the server now has to maintain only a few connections to NetScaler. NetScaler then manages persistent connections to multiple clients in a process similar to the “multiplexing” technology that cable providers use to deliver content to multiple cable subscribers. It dramatically

reduces the amount of work the server is required to perform, improving server utilization and enabling a single server to support many more users.

The numbers Citrix reports from beta implementations are impressive. While a typical server can manage a few hundred to a few hundred thousand concurrent user connections, NetScaler MPX 17000 can service upwards of two million similar connections. In customer and internal tests, Web applications that required 200 servers without Citrix NetScaler 9 could perform equally well running on 10-20 servers. Obviously, businesses stand to reap significant savings in server reduction and all the cost savings that go along with it, including acquisition, licensing, support, power, etc. According to Citrix, the bottom line is that customers can potentially reduce the number of servers by up to 10 times.

Citrix also plans to package scripts to utilize Citrix NetScaler Web 2.0 Push with a variety of application server platforms.

Key Implications for IT

- *Available April 14, 2009:* This enhancement to NetScaler 9 will be available in mid-April. For existing customers with valid maintenance, it requires an upgrade only.
- *Simplifies capacity planning:* This technology breaks the link between adding users and adding servers. Since the IT organization will very quickly be able to assess how many additional users can be supported with NetScaler acting as a proxy, the addition of 10 or 100 or even 1000 new users should have minimal impact on application servers.
- *ROI:* ROI will vary from company to company, based on the actual RIAs deployed. As always, ROI calculations will require quantification of NetScaler impact on the Total Cost of Ownership (TCO) of servers and related data center savings
- *Partners advantages:* End-user IT organizations tell us that due to the complexity of many of the integrations and composite applications supported by application servers (such as WebSphere, WebLogic, and Microsoft Windows Server), performance can be an issue. Selling Citrix NetScaler 9 along with application servers will benefit vendors and customers alike.

- *Benefits to multiple IT roles:* This announcement will be of interest to multiple roles within IT, including CIOs, Directors, IT Operations Managers, Purchasing Managers, and Application Managers and Developers. IT executives should be aware of the potential cost savings these enhancements may provide. Application Managers and Developers should note that this solution can significantly reduce the coding required to optimize connections between client and server.

EMA Perspective

It has become clear that for most application-related advancements, IT Operations pays the price. Once applications are developed and released, Development teams go on to other work. However, for data center teams, new applications are continually added and virtually nothing ever goes away. Support costs continue to rise, and the irony of the matter is that IT Operations then gets blamed for the high cost of supporting these applications. For this reason, products that promise to deliver significant cost savings on the Operations side are always welcome news.

Citrix NetScaler is part of a new wave of application delivery technology specifically designed to make Web applications more predictable, less risk-prone, less expensive to support, and better optimized for mission critical roles.

Citrix NetScaler is part of a new wave of application delivery technology specifically designed to make Web applications more predictable, less risk-prone, less expensive to support, and better optimized for mission critical roles. Load balancing and data compression functions are probably the best known of these technologies, but the good news is that vendors are delivering solutions almost as quickly as devel-

opers and Independent Software Vendors (ISVs) deliver new “problems.” It is likely that connection multiplexing, similar to what Citrix is announcing, will become a “must have” addition to other vendor’s solutions, just as load balancing became an industry standard once traditional Web applications became commonplace.

By necessity, the key vendors in this space—IBM (DataPower), F5 (Big IP), Citrix (NetScaler), and Cisco (Content Services Switch)—are all technology savvy and hotly competitive. Interestingly enough, of the four vendors, Citrix is the only one whose business was developed around delivering content to the user desktop. This gives Citrix a definite edge in terms of understanding the quirks of application delivery and end-users, and is likely one reason for Citrix’ foresight in developing this early entry into the Web 2.0 content delivery marketplace.

For companies running Web 2.0 deployments, including Cloud Service Providers, investments in this or similar technology would seem to be a no-brainer. The prospect of reducing the number of servers supporting Web applications by 90% is certainly a compelling value proposition in a realm where 10% reductions are seen as significant.

The total savings will, of course, vary according to the specific technologies and applications being supported. However, with the outcomes that Citrix is claiming, it would make sense for an in-house evaluation of Citrix NetScaler 9 to be almost automatic.

This is a new technology, and everything new carries some risk. However, EMA views these NetScaler 9 enhancements in a very positive light, particularly for IT organizations whose resources are being stretched thin by the “care and feeding” of an ever-growing array of servers supporting Web 2.0.