Web Acceleration 2.0: Now No One Has to Wait

Overview

After spending years and millions of dollars deploying a web-based, mission-critical application, the last thing you want to hear is that the rollout is going poorly due to performance issues. But, as organizations continue to migrate to browser-based deployments, that's often exactly what happens. While users enjoy the virtually ubiquitous access to applications—regardless of location—enabled by browser-based clients, they are dissatisfied with the performance of web-based applications versus their old fat-client counterparts.

Combining reduced performance with an increasingly distributed workforce introduces yet another set of factors that contribute to even more disappointing performance. WAN latency and errors prevent web-based applications from being delivered quickly and often result in congested WAN links. A plethora of network and application issues coalesce into a perfect storm, impeding performance and hampering the usability of multi-million dollar investments. Adding a dose of Web 2.0 technologies to the enterprise application mix (whether as part of an existing packaged application such as Microsoft Outlook Web Access or as a completely new initiative) only exacerbates the decline in performance for remote and mobile users.

If you're worried about the performance of your web applications, you're not alone. According to research by TechWeb Network Research, 75% of IT and applications managers experience anxiety related to application performance and 25% report physical aches and pains along with their unease[1]. These fears are not ungrounded, as there are a myriad of issues that negatively impact the performance of your applications. Network, client, and data center issues all play a role in how well your web applications perform. An application delivery network can act to alleviate the challenges inherent in delivering web applications over a widely variable network.

Delivery Challenges

Understanding the types of applications being delivered can greatly improve the ability to optimize and improve their overall performance by identifying the challenges particular to different types of applications.

Dynamic Applications

Dynamic applications—occasionally referred to as "data driven" applications because of their reliance on data sources such as relational database management system (RDBMS) and external content management systems—present a unique challenge to performance. Personalization, external data integration, and dynamic images combine with constantly changing corporate data and an increasing trend to consolidate data centers to hinder application performance in numerous ways.

Static Content

While most information in web applications trends toward 100 percent dynamic, there is still a great deal of information that is not truly dynamic and yet is often marked as such by default. Logos, images, and other various pieces of corporate information do not change drastically over time, but because it is included in pages and applications that mash-up information from sources

that *are* dynamic and constantly changing, static data ends up being constantly retrieved from servers that are already overburdened.

Yet the cost and effort to rewrite an application that properly marks content is often too high to justify the performance benefits that might be gained. Third-party solutions that eliminate the effort behind, and reduce the cost of, accurate content identification are often discarded in favor of proprietary technologies that distinguish between truly dynamic and static content, resulting in a stalemate among vendors.

SOA

Increasingly, service-oriented architecture (SOA) is the primary method of integration for both applications and organizations. Internally, SOA-based services are the foundation from which critical applications and business process are built while the interoperability of standards-based application messaging enables rapid connectivity between partners on both sides of the supply chain.

As the interconnectivity between applications and organizations grow, so does the impact of latency and unreliable services on mission-critical business processes.

Gartner estimates that 50 percent of new critical applications and processes will be SOA-based in 2007, and as many as 80 percent of these new initiatives will be designed around SOA by 2010[2].

Web 2.0

Web 2.0 applications combine the issues arising from constantly changing content with increasing volume and the connection management problems inherent in service-oriented and component-based architectures. While enabling syndication of content, providing personalization, and building an increasingly community-driven web presence is in high demand today, Web 2.0 applications can have a detrimental effect on the underlying infrastructure. Often, Web 2.0 applications consume more resources per user than their legacy counterparts due to the connection-heavy technologies underlying Web 2.0 applications such as Wikis, blogs, mash-ups, Rich Internet Applications (RIA), and discussion-based forums.

Evans Data Corporation's 2007 AJAX Development Survey of nearly 400 developers indicated that nearly four of five respondents are using AJAX for web applications and 30 percent of those developers say they will enable end-user mash-up development within the year[3].

These emerging applications are highly dynamic and constantly changing, creating challenges in formulating disaster recovery plans as well as performing backups and replication on a regular basis. The timing, amount of data, and bandwidth available for replication and backup efforts all affect an organization's confidence that their disaster recovery plans can effectively be carried out.

Network Bottlenecks

Network bottlenecks are often the cause of many an application's performance woes. Congestion, packet loss, and inherent bandwidth limitations can adversely affect the performance of even the most performance-focused architected application.

RTT

Round Trip Time (RTT) is the number of milliseconds it takes for one packet to travel from one location to another, and for a response to be returned. RTT is dependent on a number of factors, including the network infrastructure, network conditions, and the physical distance between the

nodes. Because of the impact of physical distance, RTT is often associated with the "speed of light" problem. The problem is that even if all other factors are perfect, a packet will still be limited by the speed of light.

The problem isn't really RTT, as that's merely a measure of how quickly a single packet can travel back and forth, but rather what *causes* a high RTT that negatively affects application performance. Combining the inherent speed of light limitations with mitigating factors such as network congestion contributes to the overall degradation of application performance over both the public Internet as well as dedicated lines.

Packet Loss

Packet loss occurs when one or more packets of data fail to reach their destination. Packet loss can occur because of routing failures, congestion, or physical conditions such as signal degradation or faulty routers.

While packet loss is most disruptive to unreliable transport protocols such as UDP, it also has a negative impact on the performance of applications delivered over reliable transport protocols such as TCP. Applications transported via UDP are designed to handle intermittent packet loss; applications transported via TCP are not. These applications are designed with the assumption that every packet will be delivered, and in the right order.

In the event of packet loss, the TCP receiver requests retransmission of the lost packet or in some cases the sender automatically resends any packets that have not been acknowledged. While this retransmission generally means that the application will receive all its data in the right order, it causes the application to perform more slowly as TCP adjusts its expectations to reduce the number of retransmissions necessary.

Excessive retransmissions can cause congestion and fill up queues on routers, increasing the overall time to exchange data between clients and servers. Similarly, the auto-adjustment mechanisms of TCP can cause unreliable performance of applications as it slows down and speeds up based on the current conditions of the network.

Bandwidth Limitations

The most obvious network challenge to delivering secure, fast, and available web applications is the limitation of available bandwidth. Whether the application is traversing a WAN or the public Internet, bandwidth limitations will eventually become a factor in the performance of your web applications.

Not only does the capacity of any given link in the network have an effect on the speed at which packets travel, but the increasing size and number of packets required for web applications today requires more bandwidth than ever before. On networks with limited capacity this increase can cause congestion and degrade network conditions resulting in a negative impact on the performance of all applications traversing that network.

Client Bottlenecks

While the network is often cited as the culprit for most application performance woes, there are restrictions and limitations imposed by the client that also play a part in determining how well an application performs. These bottlenecks on the client affect the performance of applications in general over all networks, but contribute even more to the overall equation that determines application performance over the WAN.

Restrictive Connection Limits

In recent years, the need to limit the resources consumed by any single end-user in order to maintain availability and performance has resulted in the imposition of connection limitations on the browser. These limitations restrict the number of concurrent connections a client is allowed to make to the host, thus ensuring that no single client can consume all the resources available and overload the server. In Internet Explorer this limitation is set at 2 connections, while in Mozilla Firefox it is set at 8 per tab with a maximum of 24.

This may at first glance seem irrelevant, but considering the number of objects that generally comprise a single page it quickly becomes apparent that not only does this limitation require longer-lived connections with the server in order to retrieve each of those objects, but the synchronous nature of HTTP requests also means that it will take much longer to retrieve the entire page. The ability to retrieve objects in parallel using the browser is limited by these artificial limitations.

Web 2.0 applications are quite often comprised of many more objects than traditional web pages, and the growing use of AJAX and other asynchronous communication mechanisms are also often affected by these limits. This can cause time-outs, excessively long sessions over existing connections, and actually prevent other pages or applications from loading and performing as expected.

Inefficient Caching Mechanisms

Another bottleneck that makes itself obvious on the client is browser caching inefficiencies. Web browsers generally adhere to the HTTP standards regarding caching directives and freshness of data. When these options are correctly configured and used within applications, the browser caching capabilities actually do improve overall application performance.

Unfortunately, most applications and web servers are not optimally configured for caching and do not take advantage of the ability to mark content as suitable for caching. This results in the ability of the browser to cache content that has no positive impact on the performance of the application.

Location-based Limitations

It has become popular to deploy WAN optimization solutions to solve many of the issues associated with delivering web applications over a wide area network. These solutions work well—as long as the client is at a remote office served by a WAN optimization controller.

Unfortunately, those clients not fortunate enough to be served by a WAN optimization controller (WOC) due to location, cost, or other factors suffer from poor application performance that may inhibit their productivity and even reduce the adoption of newly deployed web-based applications.

Also suffering from the limitations of a purely symmetrical WAN optimization solution are mobile users—those who often enjoy the benefits offered by a WOC when in the office but who suffer the same pains when traveling as their non-WOC enabled counterparts.

Counting on a WOC to solve application performance issues over the WAN or the Internet also ignores the vast majority of customers and partners who are not served by a WOC and are painfully aware of every performance problems with web-based applications.

Data Center Bottlenecks

Although it's often not addressed, there are also issues and vulnerabilities that adversely affect application performance closer to home, in the data center. Luckily, issues inside the data center are often easily identifiable and unlike emerging issues with browsers, data center problems have well-known, proven solutions.

Server Overload

Experts estimate that nearly 30 percent of a server's resources—memory and CPU—are required to process SSL. That's in addition to the resources required to open, maintain, and close high volumes of the TCP/IP connections required to respond to client requests and deliver application data. If that application happens to communicate via XML, another 5 to 10 times the resources normally required to process data and execute business logic are required.

Even over the LAN—where the speed at which the client can accept data from the server rarely has an impact on performance—the overhead of SSL, TCP/IP connection management, and the processing of XML for just a few users can quickly consume most, if not all, of a server's resources.

As resources on the server dwindle, applications become slower, tying up memory and CPU cycles more, which in turn affects the server's capacity to accept new, and process existing, requests from clients. It's a Catch-22 that often leads to unacceptable response times even over the LAN and, over the WAN, can lead to applications that simply don't work at all.

Connection Management

As connection-oriented architectures such as SOA and Web 2.0 continue to become the basis upon which applications are developed and deployed, computationally intense demands of connection management continue to take their toll on servers in the corporate data center.

In the past, web applications required just a few connections and could be tied directly to the time at which specific pages were loaded. Today's distributed, loosely coupled architectures allow for ongoing and sporadic bursts of connections between the browser and the client, and often times require additional connections between servers in the data center and with applications outside the data center, such as those hosted by partners or Software as a Service (SaaS) solutions.

Servers have a limited capacity in terms of the total number of connections that can be handled while maintaining an acceptable level of performance. Network conditions, user concurrency, and frequency of requests all impact the total capacity of a server, and that capacity can fluctuate throughout the day based on often uncontrollable and unpredictable variables.

Strategic Solutions for Application Delivery

There are several strategies for solving the issues inherent in delivering web-based applications. All have strengths and weaknesses in effectively addressing the challenges inherent in delivering business-critical web applications. Each solution needs to be considered in light of how effective it will be in meeting the needs of the business to deploy business-critical applications over the Internet and the WAN.

The most effective application delivery strategy will address the key challenges and provide for the following requirements in a single platform:

1. LAN-like performance
2. High availability and reliability
3. Predictability of performance

4. Security
5. Deployment flexibility

Data Center Expansion

The typical response to reaching capacity in the data center has long been to increase the amount of processing power available by adding servers until application performance improves to an acceptable level. This has traditionally been accompanied by the distribution of data centers in an attempt to move content and applications physically closer to the end-user, thus decreasing the impact of distance.

This “expand and distribute” strategy, however, has been increasingly rebuffed as the solution to application delivery challenges due to the cost not only in tangible assets, but in long term maintenance. While this strategy does eliminate many of the network challenges associated with delivering applications over the Internet and WAN, it also increases power, cooling, and system administration costs for the additional servers and new data centers, as well as introduce new compliance risks.

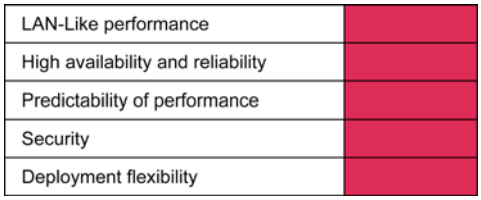

Content Delivery Networks

Content Delivery Networks (CDNs) are similar to a distributed data center with the exception that the nodes and data centers through which applications and content is delivered is owned and managed by a third party. A CDN caches static content at the edge of the network, as close to the end user as possible, thus improving performance and delivering LAN-like speed.

Unfortunately, a CDN does nothing to improve the performance of dynamic applications, nor is it designed to assist in improving the unique needs of today’s emerging architectures such as SOA and Web 2.0. CDNs are exceptionally well suited to offload static content such as images, documents, or text-based data that changes infrequently, but exceedingly poor at improving the delivery of dynamic applications or managing content that is constantly changing such as that generated by Web 2.0.

A CDN is generally controlled by a service provider who determines where nodes and distribution points will be located throughout the world. This limitation means that there are likely always end users that will not benefit from the use of a CDN unless additional nodes are deployed.

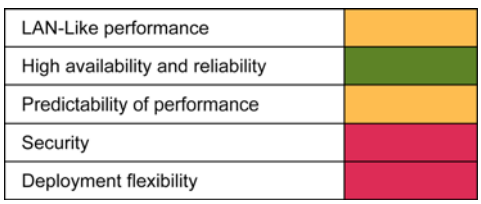

Application Acceleration Devices

Application acceleration devices have emerged to address inefficiencies in data center expansions and to optimize application delivery while providing scalability.

Application acceleration devices generally provide support for optimization of web-based applications and include the ability to compress, cache both static and dynamic application data, and offer some measure of relief for overburdened servers through connection management inside the data center.

Traditional application acceleration devices are asymmetric in their deployment, meaning that applications whose performance benefits from bi-directional optimization such as SAP, Microsoft Outlook Web Access, and Oracle are not well served by these devices.

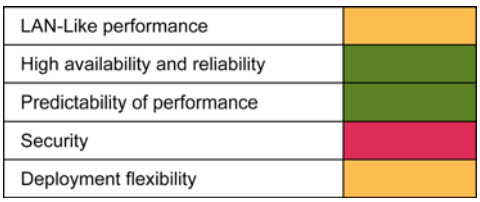

WAN Optimization Controllers

In contrast, WAN optimization controllers (WOC) are symmetric in deployment, with one device residing on both ends of a WAN link and providing rate shaping, compression, and advanced proprietary data reduction technologies to reduce the amount of data being transferred and thus improving performance.

WOCs primarily optimize content by detecting and preventing the retransmission of repeated content originating from applications within corporate data centers. This means that content incorporated from partners or third-party sites via syndication will not necessarily benefit from the deployment of a WOC.

These devices are acceptable solutions for addressing delivery challenges to remote offices but cannot offer any mitigation of performance problems for those end-users not located at WOC-enabled office. WOCs focus on solving the problems inherent in the network and do not address client or data center challenges.

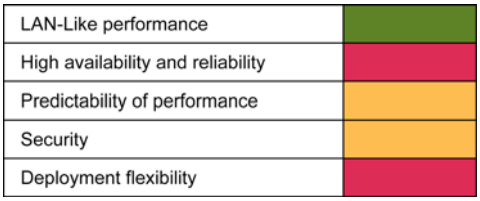

Application Delivery Networks

An Application Delivery Network (ADN) encompasses a variety of optimization, acceleration, and security techniques within a single “network” solution. This solution can be a single device or multiple devices, but regardless of the physical makeup of an application delivery network, it provides multiple features and functionality designed to address all three trouble areas of application performance: network, client, and server.

ADN provides an overlay network, positioned between the network and application layers designed to mediate between clients and applications to provide LAN-like performance, high availability, and scalability without compromising enterprise security and without requiring changes to the underlying network or the applications it delivers.

Not all application delivery networks are the same, so it is important to consider a number of criteria when evaluating an application delivery network solution:

- A purpose-built, flexible platform that addresses the inefficiencies in existing networking and application protocols to optimize application traffic
- A complete solution that addresses the bottlenecks that exist in the network, at the client, and in the data center
- Flexible deployment options that make it possible to architect the solution to meet the delivery needs of the widest possible number of end users, regardless of location
- Enterprise security should not be compromised by the solution, and in fact should be enhanced by the solution through support for security from the network to the application layer

Predictable application performance and enforcement through the implementation of adaptable service level agreements

LAN-Like performance	
High availability and reliability	
Predictability of performance	
Security	
Deployment flexibility	

F5's Approach to Application Delivery and Acceleration

F5's Application Delivery Networking portfolio is designed to address inefficiencies across the client, network, and data center while supporting a highly scalable, secure architecture. At the core of F5's integrated application delivery network are BIG-IP® solutions, operating on a modular application delivery platform. These devices address inefficiencies in the network and data center via a purpose-built operating system, TMOS™, which provides seamless integration for all F5's BIG-IP solutions.

F5's solutions comprise a breadth of technologies that accelerate, optimize, and secure applications regardless of where the end-user is located. These technologies can be deployed as individual products or integrated into a broader solution-focused architecture that is tailored to meet specific business and application needs. When deployed together, these solutions provide applications with comprehensive security, availability, and performance enhancements.

Specializing in the acceleration and optimization of web applications is F5's BIG-IP WebAccelerator™. WebAccelerator brings the next generation of acceleration and optimization techniques to F5's application delivery portfolio, offering more flexible deployment options, better security and support for security-focused solutions, and enhanced integration with F5's core platform.

Acceleration of SSL and SSL Encryption

Offloading computationally expensive operations such as SSL encryption and decryption to specialized hardware has long been an accepted method of improving performance. But, although accepted, this process brings its own issues and raises others with respect to accelerating encrypted content.

Traditional application acceleration technologies such as differential compression and byte caching do not work on SSL encrypted data. Symmetric acceleration solutions can use SSL "remoting" to address the issue, but this compromises enterprise security by introducing a convenient avenue through which DoS attacks can be launched.

WebAccelerator's solution is to intercept and decrypt the SSL data, perform the appropriate acceleration and optimization techniques on the data, and then re-encrypt it before transmitting to the receiving device. This is computationally intensive and complex, making it difficult to scale and requiring significant SSL capabilities such as cipher validation, FIPS certification, client certificate support, and hardware offload capabilities.

WebAccelerator is the first symmetric-capable application acceleration solution to include acceleration of content via SSL with hardware offload, FIPS certification, client certificate support, and other advanced SSL features.

Intelligent Browser Referencing

Intelligent Browser Referencing (IBR) technology is WebAccelerator's unique set of features that address the inefficiencies inherent in the client. IBR eliminates the need for the browser to download repetitive or duplicate data, and ensures the best use of bandwidth by controlling browser behavior. By reducing the extra conditional requests and excess data (re)transmitted between the browser and the web application, IBR mitigates the effects of WAN latency, networking errors, and packet loss. IBR also significantly reduces the amount of data downloaded without requiring Java applets or making changes to the browser. This is important, because it means IBR is a truly transparent technology. By contrast, other solutions that use applets can destroy an organization's ability to troubleshoot a site.

Key Benefits

- Alleviates bottlenecks on the client that impact application performance
- Addresses connection management issues on the client and in the data center, improving application behavior efficiency

IBR includes the following capabilities:

MultiConnect enables browsers to open more simultaneous connections between the browser and web application, allowing increased parallel data transfers. It's also extremely effective on high-latency/high-bandwidth networks such as satellite and mobile networks.

Dynamic Content Control eliminates the download of repetitive data by ensuring that the browser downloads only data that is truly dynamic and unique. It also eliminates browser "conditional requests" for static data that is incorrectly identified as dynamic while ensuring truly dynamic and unique content is freshly served.

Dynamic Linearization serves up individual pages of Adobe PDF documents from large nonlinear PDFs, allowing fast first-page views of PDF documents. Only the pages that a user is reading are transferred; users no longer have to wait for an entire manual, customer form, design spec, or drawing to be loaded prior to viewing.

Dynamic Caching

Dynamic Caching completely changes the caching model, making it possible to cache a much broader variety of content including highly dynamic web pages, query responses, and XML objects.

This patented technology is completely unique to F5 and unavailable from any other vendor. Dynamic Caching focuses on application logic and behavior, not just individual web objects. By understanding an application's high-level logic (what can and cannot be cached, what events cause invalidation, and so on), WebAccelerator eliminates repeated processing of complex web requests. Dynamic Caching enables WebAccelerator to decide when to invalidate objects and how to identify single, and reusable, pieces of content. This is made possible by predefined application acceleration policies, an intuitive user interface, a powerful XML-based API (ESI), and an HTTP request-based triggering facility that together provide comprehensive controls for validating and invalidating content.

Key Benefits

- Reduces bandwidth usage
- Reduces burden on servers
- Reduces usage of CDN services, eliminating the variability of burst charges
- Improves application performance

Without WebAccelerator and Dynamic Caching, an existing caching solution has only the object expiration date as a guide. Dynamic Caching enables the cache to look at anything in an HTTP request—from URLs to cookies, query parameters, and other headers—and produce "smart" invalidations and cache keys.

By leveraging Dynamic Caching, WebAccelerator can respond directly to up to 80 percent of the most computationally expensive user requests without involving the rest of the site infrastructure. In addition, WebAccelerator will not be confused by application semantics, and never sends invalid items from the cache

Compression

Compression is a CPU intensive operation. The ability to offload those computations onto an application delivery network at wire speeds is another important factor to consider when selecting web acceleration solutions. By taking a document as it comes back from the server and reducing it to a smaller size, it also reduces the amount of bandwidth. But more important than bandwidth reduction is this simple fact: if you're sending less data across the wire, it takes less time to get there. From that standpoint, compression is a win-win. It works both in the server offload, and also in the network and application offload, making it so the user response time is faster by getting the data across faster because there is less to transfer.

Key Benefits

- Offloading compression frees up resources on servers
- Dynamic compression capabilities ensure that compression is used only when it will improve performance

Symmetric and Asymmetric Acceleration

Traditionally, application acceleration technologies are symmetric or asymmetric, not both. This limited flexibility in deployment options impairs the ability of the organization to realize the maximum benefits for the most end users without increasing the cost and complexity of deploying a given solution.

WebAccelerator breaks this barrier by supporting both asymmetric *and* symmetric deployment models simultaneously. The same WebAccelerator device supports symmetric acceleration of web applications to remote offices while providing asymmetric acceleration to end-users accessing applications from outside remote office locations.

In a symmetric deployment, a WebAccelerator device is deployed at a central location as well as at one or more remote locations to accelerate end users' access to a web application or content that is served from the central location. The additional acceleration gains for a web application in a symmetric deployment can range from 200 to 1000 percent.

Key Benefits

- Symmetric deployment option dramatically improves application performance for users located at remote offices
- Support for simultaneous asymmetric *and* symmetric acceleration provides accelerated performance for all users regardless of location

These gains are derived from a number of functionalities working in concert. The remote WebAccelerator device acts as an edge cache for web applications and content, caching static and dynamic web objects and serving them to repeat users over the faster LAN. This removes the requirement that client web requests and responses must traverse the slow and high-latency WAN links for commonly shared objects in web applications. This is especially valuable for users that are connected via satellite or mobile links.

The nature of web sites allows for certain content to be pre-positioned in remote locations, so that those remote users have fast, local access to that content. This is useful in anticipating "flash crowds" for new content such as e-learning, online news, new products, events, and promotions. In addition to edge caching features, optimized compression, and TCP and HTTP protocol optimization is utilized between each WebAccelerator to further accelerate content, resulting in larger acceleration gains than can be achieved in an asymmetric deployment.

Proven Results with Positive Benefits

The importance of application performance on the business and end users is such that claims of improving performance should always be validated. While it's impossible to validate the improvement in performance a specific organization will gain without actually interacting within the environment, customer case studies, internal testing, and application partner testing and validation can provide at least a minimal foundation for determining the potential for performance improvements within your organization.

F5 has extensive partnerships within the enterprise application space that go beyond a traditional "works well together" partnership. F5 and its application partners work together to test, tune, and validate specialized acceleration policies across its entire application delivery networking portfolio, and specifically delivers pre-defined acceleration policies with WebAccelerator for the most popular web application platforms. These validated application acceleration policies allow you to quickly configure and deploy F5 web acceleration products to accelerate your web applications.

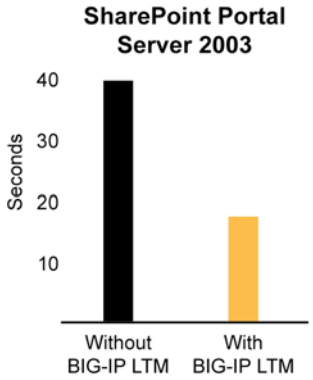

BIG-IP optimization features enhance performance of SharePoint Portal Server deployments 125% in testing with the Gomez Performance Network

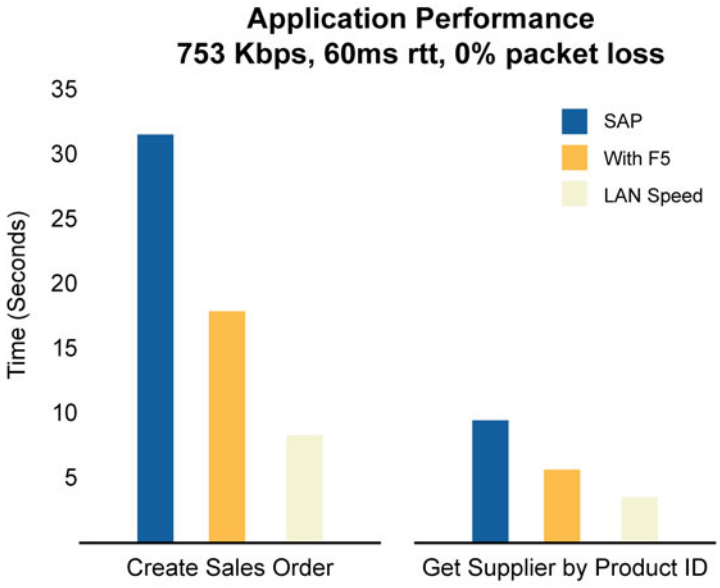

Application Performance improvements with F5 from the 768Kbps test case. Performance gains for document download on the low speed 768 Kbps network were substantial. Performance gains came primarily from the advanced compression capabilities and TCP optimizations in Local Traffic Manager, WebAccelerator, and WANJet.

Validated web application acceleration policies include:

- Microsoft Sharepoint
- Oracle Portal
- Microsoft Outlook Web Access
- Siebel 7.7
- Plumtree (BEA Aqualogic)
- Hyperion Financial

These policies can be used as templates to create custom policies for your own custom or homegrown implementations and web applications. Additionally, the policies contain in-depth knowledge of the interoperability issues between specific browsers, web applications, and network environments to help you avoid problems specific to your environment.

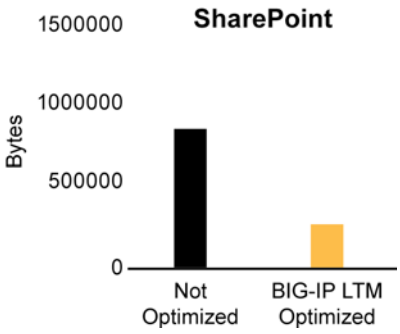

Sample Bandwidth Reduction using BIG-IP's LTM's Intelligent Compression and TCP optimizations with SharePoint 2003

Business Impact

F5's application delivery platform enables secure and fast access to applications while ensuring availability and reliability. This is beneficial to IT—especially the 75 percent feeling the effects of application-related anxiety—and also to the business overall.

Business benefits of WebAccelerator include:

- Improved application performance across the WAN, which accelerates adoption of new applications and increases end-user productivity
- Decreased variable burst costs associated with CDN services
- Reduced ongoing maintenance costs due to flexible deployment options, helping to maintain higher end-user satisfaction with application performance
- Accelerated SSL encrypted application traffic, which improves security while increasing application performance
- Reduced power and cooling costs through improved capacity of existing servers, helping to eliminate the need to expand data centers

Summary

The demand for (and benefits of) web-based interfaces combined with an increasingly mobile and distributed workforce has exposed inefficiencies inherent in the browser, the network and its protocols, and the data center. These inefficiencies create performance problems for applications that continue to be magnified by emerging technologies like Web 2.0 and SOA. The ability to scale applications, as well as networks, to provide a high level of application performance—regardless of end-user location—has become paramount to business success in today's hyperconnected world.

There is no single source for application performance woes, no single protocol or device that can be pinpointed and "fixed" to improve application performance and scalability. The unique combination of network, client, and application-specific variables all affect the overall performance of any application. All three bottlenecks must be addressed in order to improve performance, ensure availability, and enhance the security of mission-critical applications.

Moreover, these issues must be addressed with flexible deployment options that can easily adapted into unique application delivery environment. F5's approach to Application Delivery Networking eliminates the bottlenecks inherent in browsers, in the network, and in the data

center. WebAccelerator is specifically designed to enhance and improve application performance in myriad deployment scenarios; in conjunction with the rest of F5's application delivery portfolio, it helps provide a comprehensive delivery network that ensures all your applications are fast, secure, and available. For more information on F5's Application Delivery Network and its portfolio of delivery solutions, visit www.f5.com.

[1] *InformationWeek*, "Tech Pros Are Worrywarts, Says Study," March 2007 http://www.informationweek.com/story/showArticle.jhtml?articleID=197801343

[2] "Gartner's Position on the Five Hottest Topics and Trends in 2005," May 2005 http://www.gartner.com/DisplayDocument?doc_cd=125868

[3] 2007 AJAX Development Survey, Evans Data Corporation, April 2007 http://www.evansdata.com/press/viewRelease.php?pressID=3