Deploying F5 Networks BIG-IP Local Traffic Manager with WebAccelerator for Microsoft Office SharePoint Server 2007

Introduction

Organizations are using Microsoft® Office SharePoint® Server 2007 to deploy intelligent portals that provide connectivity between people, processes, customers, and intellectual capital. Specific examples include business process solutions such as team collaboration, data records management, document management and publishing, and dynamic score-carding. The growth rate is increasing, with adoption of Office SharePoint Server 2007 which provides increased functionality at somewhat higher resource costs (server CPU, memory, network, etc.) than SharePoint 2003. Many organizations who are moving to the new version are also taking the opportunity to consolidate their distributed deployments and SharePoint farms into fewer centralized data centers. This approach has many proven cost benefits, but requires these new data centers be able to provide consistent, predictable, and adequate service to users who are now located more remotely from the SharePoint farms across a corporate Wide Area Network (WAN).

In distributed farm deployments, server capacity (usually the Web front-end (WFE) servers) was likely the limiting resource to available farm capacity, with local users being connected over the corporate Local Area Network (LAN) at backbone bandwidths of 100Mbit/sec or even 1Gbit/sec. By moving the farms to a centralized location, potential limits to capacity and performance based on available WAN bandwidth (for example, a T3 connection at 44Mbit/sec, or even a T1 connection at 1536Kbit/sec) can be created. Further, the introduction of latency (delay) due mainly to connection distance, and some small percentage of dropped packets will have an impact on performance.

Why has the network now become more important?

Figure 1 shows example network traffic between end users, a SharePoint WFE server, and SharePoint database server (Microsoft SQL Server). The network traffic data observed in the HP Solution Alliances Engineering (SAE) Nashua Labs was running at a throughput rate of about 20 requests per second (RPS). Each such request represents a user 'clicking' on a portal page to perform a function (examples: home page, search, check-in document, download document, open a document library, list, etc.). Thus dividing the KB/sec values in Figure 1 by 20 generalizes the average network resource cost (traffic rate) of a single user function. The values will vary based on exact workload functions and mixes, document sizes, etc., but are a representative example.

Figure 1. Example Network Traffic rates at 20RPS throughput

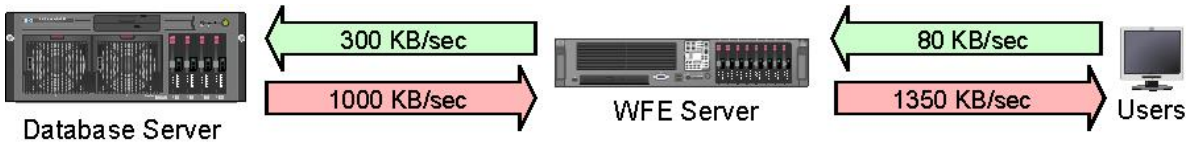

The dual-core HP c-Class BladeSystem WFE servers tested (see section *Example solution configurations – HP servers*) can each achieve about 80RPS (running WFE and Query Search services), and tested quad-core servers about 160RPS per WFE. Thus the network traffic between the WFE and users at a single server capacity of 80RPS can be more than

5000KB/sec, or about 40Mbps. This traffic, produced by a single powerful WFE server, thus exceeds the usable capacity of a T3 WAN (44Mbps). As the WAN average bandwidth usage exceeds about 75% and approaches saturation, user response times will increase dramatically, throughput will fall, users typically will become frustrated, and the solution becomes inefficient. Technologies to reduce (optimize) WAN bandwidth use and minimize the effects of latency are thus highly desirable.

One approach to improving overall system performance is the use of a class of Traffic Management Device commonly referred to as a WAN Accelerator. This paper details the results of using one such solution – the F5 Networks® BIG-IP® Local Traffic Manager™ (LTM™). By enabling the WebAccelerator™ module on BIG-IP LTM, test results showed a significant increase in user-perceived performance while at the same time decreasing the load on the WFEs. This technology was tested in the HP SAE Nashua Labs in two typical customer scenarios. A "branch office" scenario representing corporate users accessing SharePoint performing read/write operations; and an "Internet" scenario representing anonymous users accessing (read only) a company's externally-facing Internet portal. The two solution deployments, best-practice configurations and observed performance results are detailed in the following sections; and demonstrate how significant performance improvements can be realized.

NOTE: Throughout this document, all references to BIG-IP refer to BIG-IP Local Traffic Manager (LTM) with the WebAccelerator module enabled.

Audience

This paper is intended for people who will be proposing solutions, providing installation services or consulting, and who may be assisting in deploying Office SharePoint Server 2007 solutions accessed over a Wide Area Network (WAN). It will also be of interest to IT professionals who may be deploying and/or managing such Office SharePoint Server 2007 solutions.

HP recommends the information provided herein be used in conjunction with the Office SharePoint Server 2007 product solution documentation; and information contained in additional white papers and articles authored by the WAN technology solution vendor, F5 Networks Inc., as noted in the section titled *For more information*.

Solution overview

The solutions described employ F5 BIG-IP LTM with WebAccelerator appliances in both symmetric and asymmetric deployments, as applicable for the two example scenarios (branch office and Internet).

For this test engagement, HP SAE Labs selected to use BIG-IP LTM with the WebAccelerator software module. BIG-IP LTM provides the load balancing functionality to ensure high availability and scalability, while WebAccelerator focuses on increasing overall HTTP and HTTPS performance. Although the two products can be purchased as separate devices, the combined functionality provides a best-of-breed solution at a cost effective price.

The F5 BIG-IP LTM with WebAccelerator platform selected was the 6800, which is a 2U, port dense network appliance capable of handling nearly 4 Gbps of sustained throughput.

The hardware runs F5's TMOS™ (Traffic Management Operating System) architecture platform, which was built from the ground up as an OS whose sole purpose was application delivery networking. The end result is a platform that delivers unparalleled data inspection, manipulation, and acceleration at line speed. With more than 16,000 customers, F5 has more than 60% of the worldwide market share of the advanced platform Application Delivery Controller market for 2Q07 based on revenue (*source: Gartner, Market Share: Application Acceleration Equipment, Worldwide, 2Q07, October 2007*).

In a symmetric deployment, two such BIG-IP LTM with WebAccelerator appliances located at each 'edge' of the WAN can work collaboratively to provide significant WAN traffic reduction and optimization. A single appliance in an asymmetric deployment can not provide all of the features of a symmetric configuration, but can still result in notable improvements in bandwidth usage and user-perceived performance.

Technologies

The BIG-IP LTM with WebAccelerator appliance provides a number of technologies and features that can be enabled as part of a SharePoint deployment. These technologies and features are designed to consistently delivery data to the end user in the fastest method possible while also decreasing server and bandwidth loads. They include:

- Hardware assisted network routing and bridging
- Traffic management (for example, WFE server load balancing)
  - Application health monitoring
  - Server performance based load balancing algorithms
- HTTP/S application acceleration technologies
  - Dynamic end point caching
  - Intelligent compression
  - SSL acceleration
  - TCP multiplexing
  - Intelligent Browser Referencing (IBR)
  - SharePoint layer 7 (object) protocol optimization
- TCP optimizations (optimizes connections for both local and remote users)
  - Delayed and selective acknowledgements
  - Explicit congestion notification
  - TCP slow start
  - Congestion avoidance
  - TCP Spooling

These capabilities can be selected as appropriate in each scenario to maximize the efficiency of the WAN. The best practice for configuring these options is to utilize the BIG-IP WebAccelerator "SharePoint profile" (see later section *Solution set-up*). As noted earlier, it is the traffic from the WFEs to the users' desktop browsers that comprises the vast majority of the WAN traffic. This traffic can include HTTP, Java™, XML, etc. to provide page content, and various Microsoft Office documents or other files 'opened' by the user. Reduction in that traffic by optimization, compression, and by client-end caching will improve performance in differing degrees in both example scenarios.

When tailored to a specific SharePoint deployment, the preceding technologies should provide:

High Availability – BIG-IP LTM with WebAccelerator is a critical piece of providing a highly available SharePoint infrastructure. By configuring BIG-IP LTM with WebAccelerator to monitor the WFEs in order to determine if they are serving valid content or not, users will only be sent to the servers that are truly available. A properly configured BIG-IP LTM with WebAccelerator will not send any users to a WFE that is mis-configured, down, or saturated.

By implementing the overall solution with redundancy in the BIG-IP LTM with WebAccelerator appliances, SharePoint Servers and Services, as well as a SQL clustered back-end, it is possible to build a system that can truly withstand the loss of any one component.

Scalability – As enterprises grow, it is critical that information systems can scale to serve the needs of its users. Single server deployments can scale by building higher-end systems; however there is a measurable point in which this becomes a cost ineffective model. By deploying BIG-IP LTM with WebAccelerator, it is possible to scale out multiple WFEs and spread the user load amongst the farm. By optimizing the traffic via technologies like caching and TCP multiplexing, you can actually scale the system without experiencing server proliferation or bandwidth saturation.

Performance – Countless studies have shown that user adoption of a corporate resource will suffer as user perceived performance of the system declines. In order to make sure that all users, whether local or remote, have the best overall user experience, it will be highly beneficial to leverage BIG-IP LTM with WebAccelerator SharePoint Acceleration functionality.

Branch office scenario

The branch office scenario represents a centralized data center supporting Office SharePoint Server 2007 farms, and a number of authenticated users co-located at a remote location connected by a corporate WAN of known bandwidth, latency, etc. The WAN characteristics will likely depend on geographic location and span; and the network usage will depend on the number of users at the remote location and the workload (functions and frequency) that they apply to support their business needs (SharePoint and other traffic).

A BIG-IP LTM with WebAccelerator symmetric deployment best suits this scenario. BIG-IP LTM with WebAccelerator appliances will be located both at the data center and at the remote location, and be configured to work co-operatively. This configuration can provide all the features noted above, including very high ratio compression/decompression of content by the appliances. The remote appliance can also intelligently cache content local to the users, thus greatly reducing WAN traffic. An appliance can be located at each such remote office location as needed. The number of appliances at the data center will depend on the appliance capacities (different models providing varying capacity) and on high-availability (fail-over) needs.

Internet scenario

The Internet scenario depicts a centralized data center providing Office SharePoint Server 2007 sites being accessed by anonymous users on the Internet (outside the corporate network perimeter). These users will be consumers of information (read-only access) and the corporate network will be protected by appropriate security (for example, deployment of Microsoft Internet Security and Acceleration (ISA) Server) at the network perimeter (edge of the WAN).

A symmetric configuration is not feasible in this case, thus an asymmetric single BIG-IP LTM with WebAccelerator appliance will be used at the data center, located behind the network perimeter. This configuration can not achieve all the capabilities of a symmetric BIG-IP LTM with WebAccelerator deployment, however it still provides a good degree of compression efficiency as it can use the users' browsers to decompress content; and caching efficiency is still high as the BIG-IP WebAccelerator Intelligent Browser Referencing (IBR) technology will ensure the users' browser caches are maintained current. Test results show an asymmetric deployment provides notable WAN optimization and improved performance.

Example solution configurations

The following sections provide details of the solution components and configurations as tested in the HP SAE Labs.

Software components

The following software components were utilized on the various Office SharePoint Servers in the farm, each as either X86 (32-bit) or X64 (64-bit) versions as needed:

- Windows Server 2003 SP1 (on all servers)
- Microsoft SQL Server 2005 SP2 (on database servers)
- Microsoft Office SharePoint Server 2007

Once all software was installed, the SharePoint 2007 Configuration Wizard was used to create the farm Administration service, and then the required WFE, Query Search, and Index Search services. Test farm sites were created and content deployed to document libraries, lists created (for example, events, announcements), and other site pages and content to support the standard HP test workloads (described later).

HP servers

The test farm comprised a total of seven HP c-Class blade servers representing an enterprise-class solution deployment, configured as:

- 2 x HP ProLiant BL465c, Dual-Core AMD Opteron™, 2.80GHz – WFE + Query Search
- 2 x HP ProLiant BL460c, Quad-Core Intel® Xeon®, 3.0GHz – WFE + Query Search
- 1 x HP ProLiant BL460c, Quad-Core Xeon, 3.0GHz – Index Search
- 2 x HP ProLiant BL460c, Quad-Core Xeon, 3.0GHz – Database servers (active/passive SQL cluster)

As noted, the first four servers were running both the WFE and Query Search services and were configured as a load-balanced pool to test the BIG-IP LTM with WebAccelerator load balancing capability (as opposed to using Windows Load Balancing Service). Normally servers of the same 'power' would be used, however BIG-IP LTM with WebAccelerator provides multiple load balancing methods that can effectively balance load across servers of disparate power (see the section titled *Solution setup* for configuration details).

As the throughput of a single WFE of this power can exceed the usable network capacity of a 44Mbps WAN, the WAN tests ran against a single BL460c quad-core server. This also provided a single consistent test platform for all tests.

WAN acceleration hardware

Figure 2 depicts two BIG-IP LTM 6800 with WebAccelerator appliances deployed in a symmetric configuration. Authenticated Users were connected (via a network switch) to the client-side BIG-IP on a remote LAN. The server-side BIG-IP LTM with WebAccelerator was connected to the HP c-Class enclosure internal network switch. The other ports of the BIG-IP LTM with WebAccelerator appliances were connected to the WAN (a WAN emulator was employed in the HP SAE lab tests).

Figure 2. Symmetric F5 BIG-IP LTM with WebAccelerator deployment

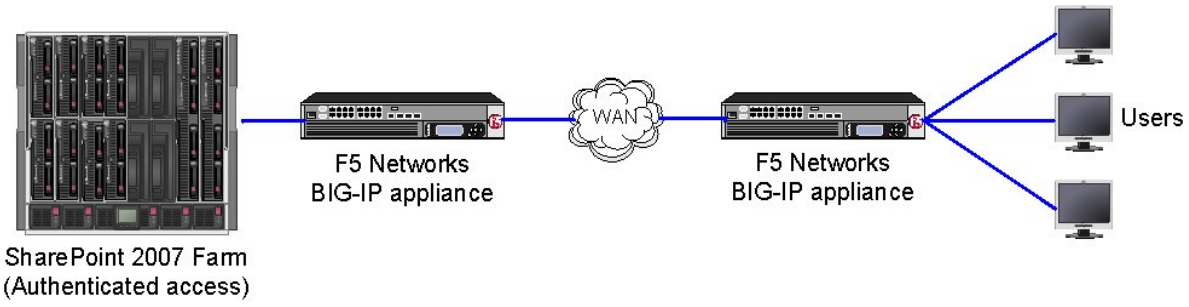

Figure 3 shows a single BIG-IP LTM 6800 with WebAccelerator appliance connected to the HP c-Class blade enclosure internal network switch. The other BIG-IP port was connected to an ISA server (providing corporate LAN security). Anonymous users connected via a network switch to the WAN, and then to the ISA server.

Figure 3. Asymmetric F5 BIG-IP LTM with WebAccelerator deployment

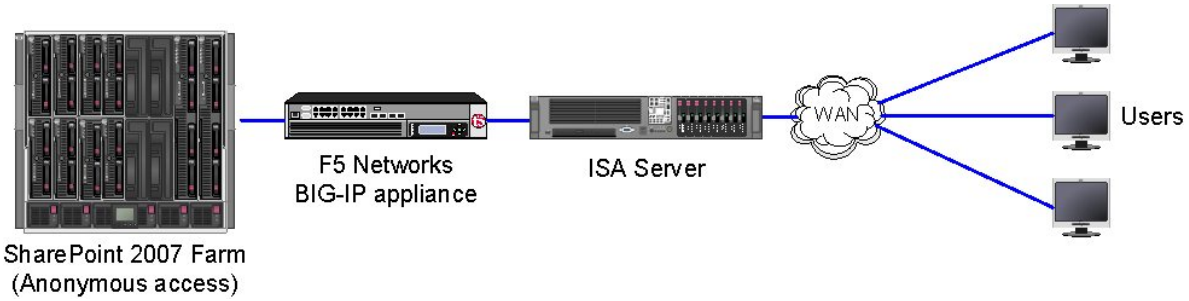

The following describes best-practice solution deployment guidelines, and BIG-IP LTM with WebAccelerator configuration details as tested in the HP SAE labs.

Solution deployment

These sections summarize key aspects of typical solution deployments, although individual customer deployment specifics may vary. For detailed information, refer to the *F5 BIG-IP SharePoint 2007 Application Ready Network Deployment Guide*, a pointer to which is provided in the section *For more information*.

Installation

Physical installation of the BIG-IP LTM with WebAccelerator appliances is simple, each requiring a 2U rack space and power to the rear-mounted power supply socket(s). In most cases, the BIG-IP LTM with WebAccelerator appliances will be implemented in a redundant fashion, with two units operating in an active/standby pair.

An 'out of band' management network port on the front of the BIG-IP LTM with WebAccelerator provides secure access to the BIG-IP subsystem. This port is connected to a dedicated onboard processor that can provide lights-out management, monitoring, and diagnostics of the BIG-IP LTM with WebAccelerator appliance. Basic management and monitoring can also be achieved via the display screen on the front of the unit.

Network cables for production traffic can be plugged into the tri-speed copper/fibre ports found on the front of the BIG-IP LTM with WebAccelerator appliance. These ports are connected via a true layer-2 switch plane, and allow for the traffic to pass through the traffic management engine as necessary.

Solution set-up

The BIG-IP LTM with WebAccelerator management port ships with a default IP address, however it is easily modified via the front panel in order to match your network. The BIG-IP administrator can use that IP address to access the comprehensive browser-based configuration page via HTTPS and a powerful command-line interface (CLI) via SSH. Further setup and configuration is performed via various hierarchical screens in the web interface.

F5 has documented the step-by-step procedures for configuring the BIG-IP LTM with WebAccelerator system to manage and accelerate traffic in a SharePoint 2007 environment in their *F5 BIG-IP SharePoint 2007 Application Ready Network Deployment Guide*, a pointer to which is included in the section *For more information*. Please refer to this document for detailed setup procedures.

The following will describe an example SharePoint 2007 and BIG-IP LTM with WebAccelerator deployment (as tested in the HP SAE Labs) by showing some key BIG-IP administration screens and explaining their purpose. Although not an exhaustive configuration guide, it will show some of the settings that should be configured in a typical deployment. The specific BIG-IP LTM with WebAccelerator features employed as best-practices in asymmetric and symmetric deployments are also discussed.

Example setup screens

Figure 4 shows the Virtual Servers Properties tab. The General Properties show a virtual server name (VS_Wan), a destination host IP address and the port (HTTP port 80) being monitored. The Configuration section shows that the underlying protocol is TCP, and that the HTTP Profile "http-wan-optimized-compression" is being used.

Figure 4. BIG-IP LTM with WebAccelerator setup – Virtual Server general properties

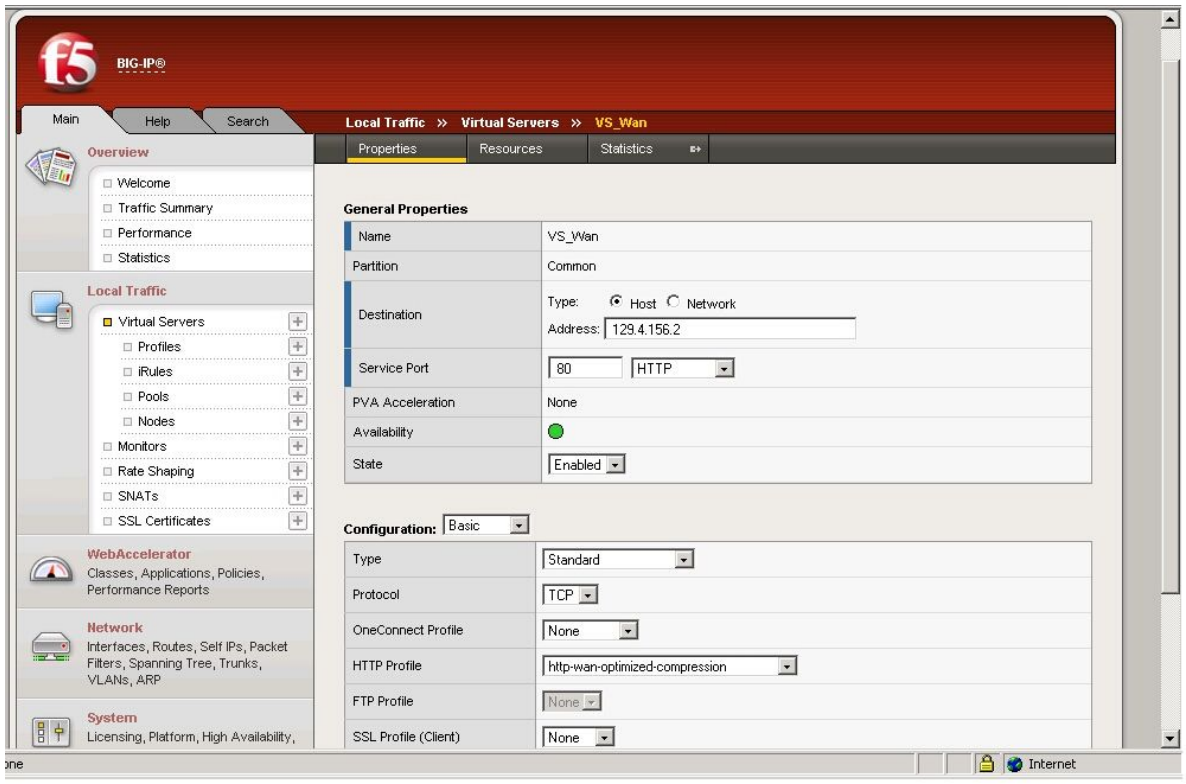

Figure 5 shows the Virtual Server Profile Properties page. This enables definition of a Class based on a number of provided profiles which can be customized. Of note is the selection in the Configuration section showing WebAccelerator is enabled (as noted above). The selection of "WebAccelerator Enabled" in the profile properties provides client-end caching functionality, thus using the HTTP profile "http-wan-optimized-compression" per Figure 4 (as opposed to "http-wan-optimized-compression-cache") is the correct selection.

Figure 5. BIG-IP LTM with WebAccelerator setup – Virtual Server profile properties

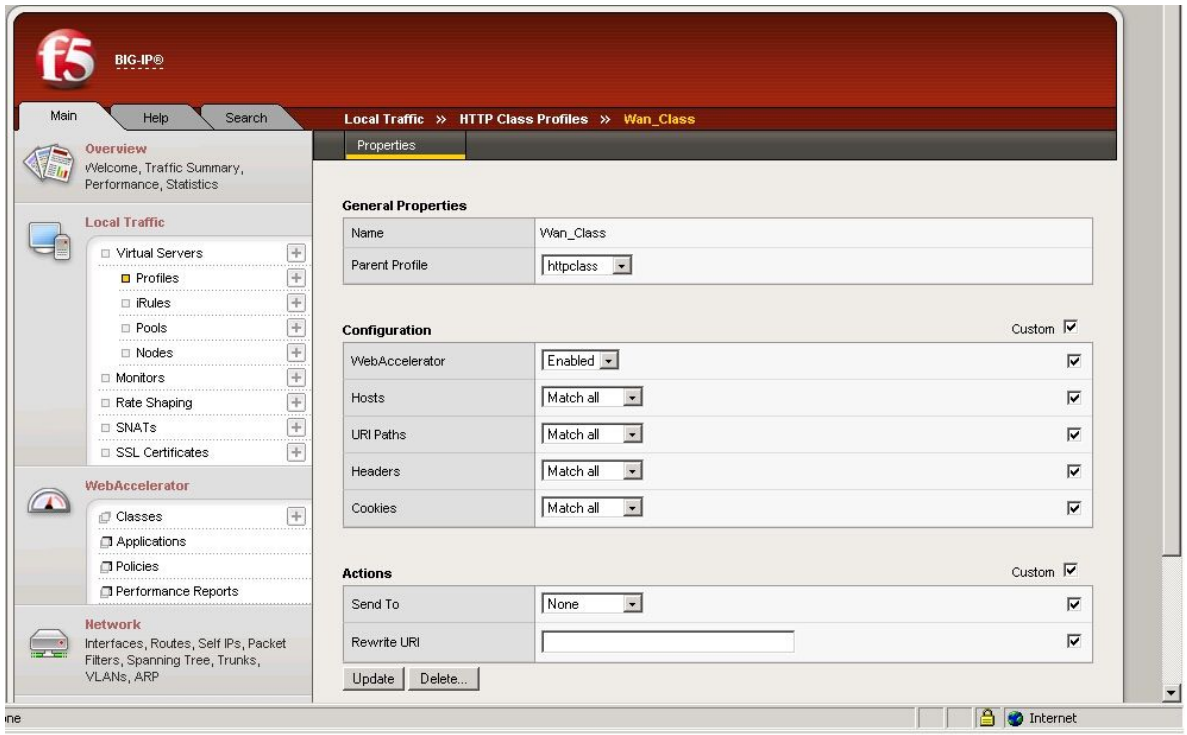

Figure 6 shows the Virtual Server resources screen. This enables selection of the load balancing pool ("SharePointFrontEnds pool" shown defined in Figure 7) and of the HTTP Class Profile ("WAN_Class", as shown defined in Figure 5 above).

Figure 6. BIG-IP LTM with WebAccelerator setup – Virtual Server resources

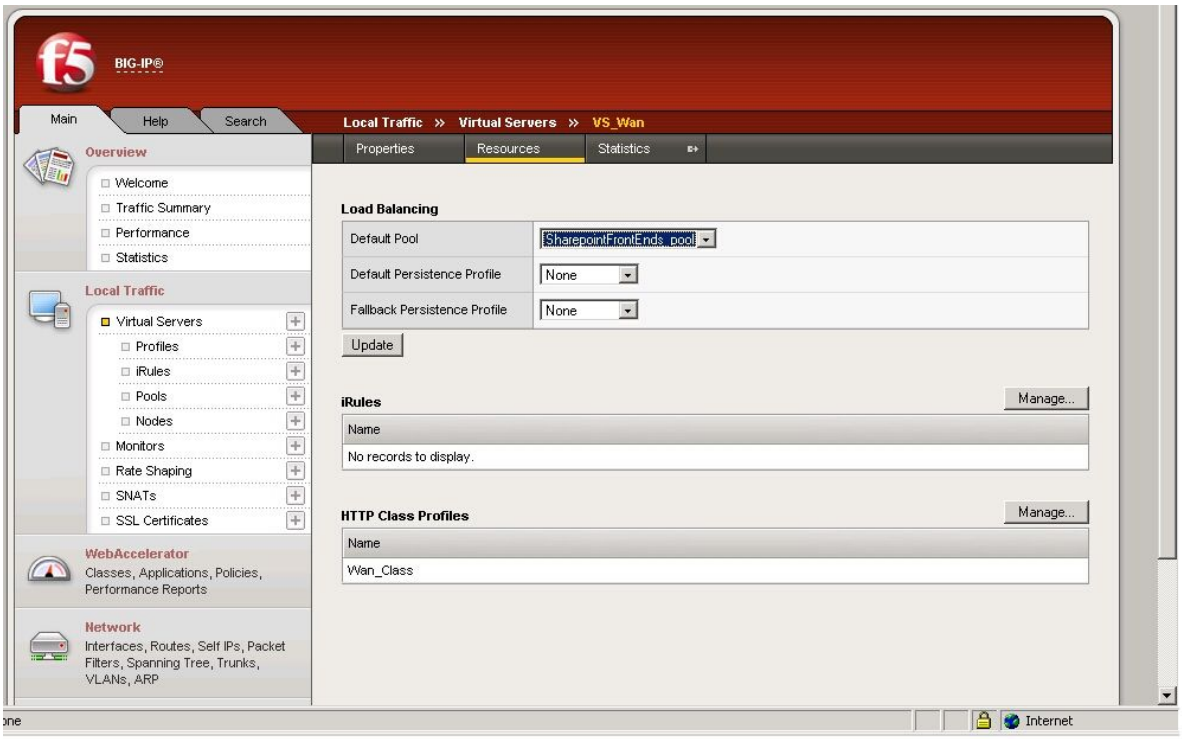

Figure 7 shows the Load Balancing pool definition as regards the pool server IP addresses and the load balancing method choices. Various server configurations and load balancing methods were successfully tested. Balancing servers of disparate power is possible by applying weighting to the servers to reflect their individual power or throughput capacity. In the example below, only one server (IP address 129.3.30.113:80) is shown with a green 'active and enabled' symbol, as subsequent WAN testing was performed using only that one server as noted earlier. The black diamond symbol denotes an 'inactive' server that could not be contacted (powered off), and the black circle symbols represent active servers (can be contacted) but are not currently members of the load balanced pool (disabled).

Figure 7. BIG-IP LTM with WebAccelerator setup – Virtual Server pool members

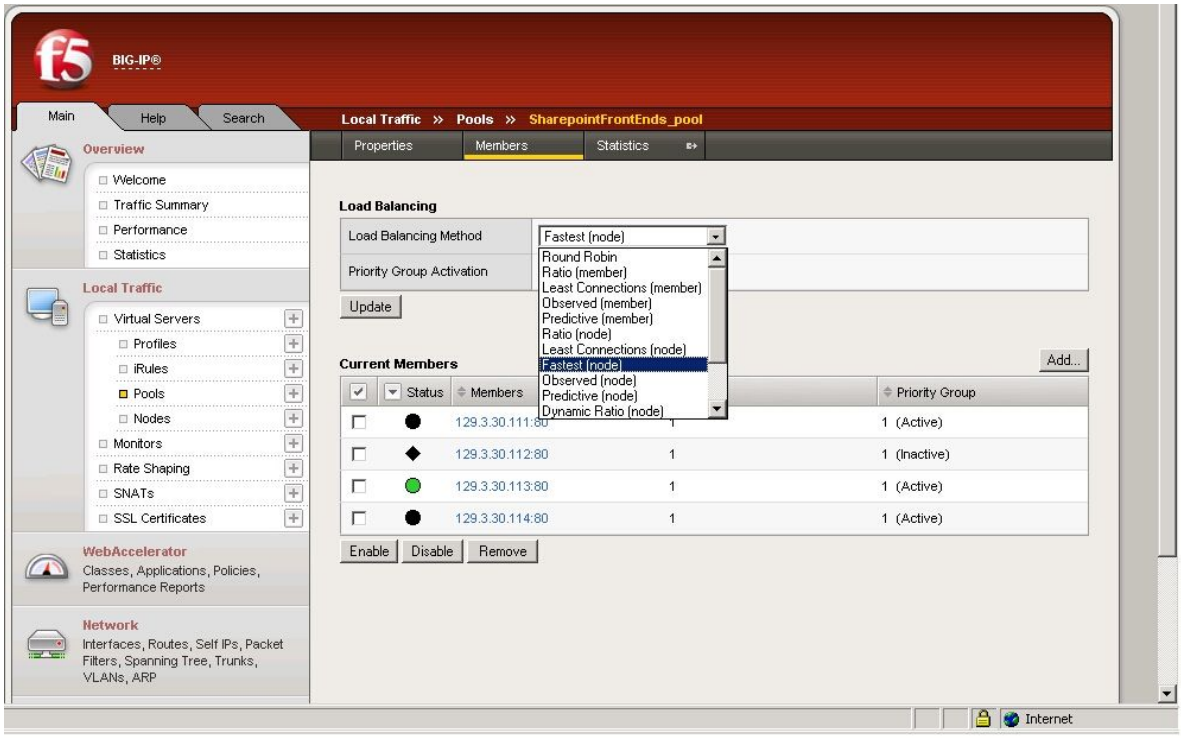

The following sections provide details regarding the BIG-IP LTM with WebAccelerator features provided with asymmetric deployment; and additional features from symmetric deployment.

Asymmetric operation

By front-ending the SharePoint WFEs with BIG-IP LTM with WebAccelerator, remote users connecting to the data center for content will notice a significant decrease in page retrieval time. There are several technologies that make this happen. These include (but are not limited to):

BIG-IP WebAccelerator Dynamic Caching – Caches unchanging data that may seem dynamic (contains query parameters, etags, session IDs), but is actually static data or changes in an identifiable pattern. BIG-IP WebAccelerator can cache a higher percentage of data from

dynamic web applications while maintaining proper application behavior. It accomplishes this by fully inspecting every aspect of HTTP requests, controlling caching behavior, and invalidating cached data.

BIG-IP LTM SSL Acceleration – Offloads the servers from computational-intensive SSL encryption and decryption, reducing server processor utilization by as much as 50%.

BIG-IP WebAccelerator Dynamic Compression – Enables BIG-IP WebAccelerator to compress dynamic data from web applications. BIG-IP WebAccelerator Dynamic Compression is different from standard compression implementations because of its compression efficiency, and its ability to avoid widespread browser compression bugs. Further enhancements are utilized when serving dynamic, unique, or modified requests for compressed data from the cache. Even dynamic content requiring unique session IDs within every link on the page can be delivered and compressed with zero compression overhead.

BIG-IP WebAccelerator Intelligent Browser Referencing (IBR) – IBR is a group of capabilities that eliminates the need for the browser to download repetitive or duplicate data, as well as 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 reduces the effects of WAN latency and errors. IBR also significantly reduces the amount of data downloaded without requiring Java applets or making changes to the browser that are common in delta compression methodologies.

IBR is comprised of three main functionalities – MultiConnect, Dynamic Content Control, and Dynamic Linearization.

- **MultiConnect** – Enables Internet Explorer to open more simultaneous connections between the browser and web application, allowing increased parallel data transfers. MultiConnect is extremely effective on high latency/high bandwidth networks such as satellite and mobile networks.
- **Dynamic Content Control (DCC)** – Eliminates the download of repetitive data by ensuring that the browser downloads only the data that is truly dynamic and unique. DCC also eliminates browser 'conditional requests' for static data that is incorrectly considered dynamic while ensuring truly dynamic and unique content is freshly served.
- **Dynamic Linearization** – Serves up individual pages of Adobe® PDF documents from large non-linear PDF files, allowing for 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 specification, or drawing to be loaded prior to viewing.

Symmetric operation

By deploying BIG-IP LTM with WebAccelerator in a symmetric fashion, users will request the SharePoint data from their local (or closest) BIG-IP. If the local BIG-IP does not have the requested data, it will pull it from the BIG-IP that is local to the SharePoint farm.

Users requesting content from their branch office BIG-IP LTM with WebAccelerator will benefit from the same technologies as for the asymmetric deployments, as well as:

SharePoint Data Proximity – Commonly requested SharePoint pages and objects will be cached at the branch office BIG-IP LTM with WebAccelerator appliances, greatly reducing the need to re-pull the data every time it is requested.

BIG-IP-to-BIG-IP secure data transfer acceleration – Since the BIG-IP LTM with WebAccelerator appliance becomes both end points of the TCP connection, it is able to leverage hardware and software that is designed to transfer data faster than normal browser to server connections.

Performance results

The following sections describe the tests performed, results observed, and also provide a summary of key findings and conclusions.

Test methodology

The methodology employed was designed to determine the performance improvements that could be realized by deploying the F5 BIG-IP LTM with WebAccelerator solution to two scenarios, compared to an un-accelerated baseline. More than 100 individual tests were performed, with comprehensive monitoring of all server and network resource utilization. The following were the key methodology steps in testing both scenarios:

1. Configure the scenario without BIG-IP LTM with WebAccelerator appliances (un-accelerated baseline configuration).
2. Apply the appropriate workload at increasing load levels until average WAN utilization reached approximately 75% of WAN bandwidth (best achievable un-accelerated result).
3. Repeat Step 2 for all five WAN 'types' (see Table 1, below) to provide all baseline data.
4. Deploy best-practice BIG-IP LTM with WebAccelerator configuration for the scenario.
5. Repeat test load achieved in Step 2 (best un-accelerated load) to determine effect of BIG-IP LTM with WebAccelerator at that load level (direct comparison result).
6. Continue increasing load until 75% WAN utilization once again achieved (best achievable accelerated result).
7. Repeat Steps 5 and 6 for all WAN 'types' to provide all 'accelerated' data.

Key performance metrics monitored included Client (users) observed KBytes/sec network traffic, WAN network traffic (Mbps or Kbps), SharePoint requests per second (RPS), HTTP Hits/sec, .NET Requests/sec, WFE CPU Utilization%, etc. These data provided a comprehensive view of 'cause and effect,' and demonstrated the effects of BIG-IP LTM with WebAccelerator on network traffic (LAN and WAN), WFE CPU, and achievable throughput (RPS).

The results presented below for both scenarios are a very small subset of all those recorded. They are intended to summarize the results of Steps 2, 5, and 6 (un-accelerated baseline, accelerated comparison baseline, and accelerated best-achievable). As an aid for your understanding and analysis of the results, a further section (*Key findings*) will also summarize the key observations and conclusions found from the results.

Workload scenarios

Two workload scenarios were used, each including a range and mix of common SharePoint functions that users in the scenarios would typically and frequently employ. These included Home page, list browsing, search, document library (doclib) browsing, opening/reading documents, document check-out/-in, etc. Office SharePoint Server 2007 is a feature-rich and highly customizable application, and thus customers' workloads may include functionality

not included in the test workloads. The workloads that were used, however, were intended to be broadly representative. In short – 'your mileage may vary.'

Branch office

The emulated branch office users are authenticated onto the SharePoint site, and are both consumers and suppliers of information (that is, read/write activities that modify data content). The following shows the function mix used. The percentages do not sum to 100, but are rather the percentage likelihood of a given function occurring for each pass of the workload 'script.' The script loops continuously during the test duration.

- Home 70%
- Read from Events calendar list 65%
- Read from Announcements list 65%
- Search (range of keywords) 25%
  - Open/read document from Search results set 10%
- Pick random Doclib folder and browse content 65%
  - Open/read document from folder 10%
- Pick specific Doclib folder and browse content 10…40%
  - Check-out/modify/check-in document from folder 10%

Internet

The emulated Internet users have anonymous access to the SharePoint site, and are consumers of information (that is, read-only). The following shows the function mix used. As no read/write activities are performed, and mix and percentages of the other functions are adjusted to suit (for example, more searching and browsing). The script loops continuously during the test duration.

- Home 70%
- Read from Events calendar list 65%
- Read from Announcements list 65%
- Search (range of keywords) 40%
  - Open/read document from Search results set 25%
- Pick random Doclib folder and browse content 65%
  - Open/read document from folder 25%

Network configurations and parameters

The solutions (both symmetric and asymmetric) were tested against a range of emulated WAN properties as regards bandwidth, latency and packet loss percentage. These were chosen to represent a range of typical WAN environments that will be seen globally. Table 1 details the tested WAN environments, their characteristics and examples of real-world application.

Table 1. WAN Network parameters

WAN Type	Bandwidth	Latency (mSec)	Packet Loss %
1. Regional connection (T3)	44 Mbps	90 ms	0.1%
2. U.S. branch office	6 Mbps	90 ms	0.1%
3. U.S. – EMEA	6 Mbps	200 ms	0.1%
4. U.S. – EMEA satellite	6 Mbps	300 ms	0.1%
5. U.S. – APAC (T1)	1536 Kbps	300 ms	1.0%

The first WAN type is a T3 connection at 44Mbps as might be found connecting corporate regional centers. Average latency, due to distance, is 90ms. A small percentage of packet loss will occur. Type 2 represents a U.S. branch office connection at 6Mbps. Type 3 is a connection between locations in the U.S. and Europe, with an increased latency due to the distance. Type 4 is similar, but represents a satellite (versus cable) connection with increased latency. Finally, type 5 represents a T1 U.S. to Asia-Pacific connection at lower bandwidth, a latency reflecting longer cable distance and a higher packet loss percentage.

Data table metrics

Tables 2 and 3 detail comparative test results and present a range of key performance metrics. The following describes the table column headings and metrics, to assist you in understanding of the data.

- **Load** – a relative measure of client load applied to the test system. Not directly related to 'supported users,' which can be better derived from 'Avg RPS' (see below).
- **WAN type** – refer to Table 1 for details of the five WAN configurations tested. The letter suffix indicates the result for 'U' (best un-accelerated result), 'D' (accelerated direct comparison result), and 'A' (best possible accelerated result). For example, comparing results 2U and 2D shows the difference observed when deploying BIG-IP LTM with WebAccelerator but running the same load (that is, reserve capacity gained) for a WAN running at 6Mbps, 90ms latency and 0.1% packet loss. Whereas, comparing 2U and 2A shows the actual gain in performance possible with BIG-IP LTM with WebAccelerator (higher load and throughput), both shown at the maximum recommended WAN utilization rate.
- **Avg RPS** – average SharePoint requests/second. A 'request' represents a user clicking on an item resulting in a new page being presented. This is the key performance throughput metric, directly related to user perceived response times. Response time is largely influenced by network latency, as WFE server capacity is more than adequate in all these tests. RPS is also broadly equivalent to HTTP pages/sec and to .NET requests/sec.
- **Avg hits/sec** – generally, an HTTP page will require multiple HTTP Hits to render. This will vary depending on page complexity (number of SharePoint web parts, etc.) Observation shows that typical SharePoint pages have an un-accelerated Hits-to-Page ratio of 3:1.

- **Avg client KB/sec** – the average Kilobytes/second (KB/sec) across all emulated clients as reported by the client emulation tools. This equates to the User LAN traffic rate.
- **Avg Page** – the average response time (seconds) of all workload function pages, as reported by the client emulation tool. Statistically an 'average of averages' it provides an indication of overall response time changes due to BIG-IP LTM with WebAccelerator.
- **Avg Search** – the average response time (seconds) for the Search function to return a results set page to the user. A range of different search keywords was used, thus the min/max response time for this varies.
- **Avg List Open** – the average response time (seconds) for the user to be presented with a doclib folder content list of documents after clicking open a doclib folder. Multiple folders with differing content were used.
- **Avg Doc Open** – the average response time when a user clicks on a document in a doclib list to open/read it. As differing sizes/types of Microsoft Office documents were used, the time will vary depending on the size, available bandwidth and latency.

Branch office results

Table 2 details results obtained for the Branch Office scenario, both un-accelerated and when employing a BIG-IP LTM with WebAccelerator symmetric deployment solution. Comparative data is shown for the un-accelerated (U), direct comparison (D), and accelerated (A) results.

Table 2. Symmetric configuration results

Load	WAN test	Avg RPS	Avg Hits/sec	Avg Client KB/s	Avg Page	Avg Search	Avg List Open	Avg Doc Open
350	1 U	40.00	121.00	3205	2.95	2.04	5.81	8.67
350	1 D	56.70	53.30	1897	0.44	0.49	2.08	0.65
600	1 A	93.70	88.20	3111	0.59	0.92	2.69	0.78
55	2 U	6.20	19.10	507	1.12	2.14	6.10	9.57
55	2 D	9.00	8.50	300	0.35	0.4-0	2.00	0.55
300	2 A	47.60	44.70	1619	0.50	0.66	2.58	0.61
70	3 U	6.20	19.10	506	1.97	4.51	9.97	17.41
70	3 D	10.70	10.00	362	0.73	0.76	3.90	0.88
300	3 A	44.90	42.20	1511	0.88	1.00	4.70	0.92

Load	WAN test	Avg RPS	Avg Hits/sec	Avg Client KB/s	Avg Page	Avg Search	Avg List Open	Avg Doc Open
80	4 U	6.10	18.80	498	2.83	5.04	13.53	24.45
80	4 D	11.60	10.90	387	1.11	1.16	5.81	1.00
300	4 A	42.50	39.80	1430	1.26	1.42	6.39	1.12
22	5 U	1.53	4.70	124	3.12	6.46	15.54	26.70
22	5 D	3.12	2.97	105	1.29	1.49	6.66	1.03
100	5 A	13.60	12.80	462	1.53	1.72	7.67	1.19

Internet results

Table 3 details results obtained for the Internet scenario, both un-accelerated and when employing a BIG-IP LTM with WebAccelerator asymmetric deployment solution. Comparative data is shown for the un-accelerated (U) and accelerated (A) results.

Table 3. Asymmetric configuration results

Load	WAN test	Avg RPS	Avg Hits/sec	Avg Client KB/s	Avg Page	Avg Search	Avg List Open	Avg Doc Open
400	1 U	43.40	127.40	3722	3.15	2.10	5.54	10.82
400	1 A	56.40	65.10	2523	1.04	1.20	3.32	5.55
55	2 U	4.91	14.30	417	5.16	3.39	9.61	19.42
55	2 A	7.83	9.01	356	0.96	0.97	3.18	6.06
70	3 U	4.88	14.20	415	8.25	5.19	14.86	32.01
70	3 A	8.82	10.10	394	1.90	1.89	6.28	12.40
80	4 U	4.70	13.40	409	10.90	7.31	18.92	48.20
80	4 A	9.20	10.50	402	2.71	2.58	8.78	18.07
24	5 U	1.29	3.65	108	12.70	8.24	22.53	54.10
24	5 A	2.36	2.73	97	4.16	4.60	12.45	21.90

Key findings

The following are key observations from the data presented.

Branch Office – symmetric BIG-IP LTM with WebAccelerator solution

- Results for the 6Mbps and 1536Kbps tests show that deploying BIG-IP LTM with WebAccelerator provided almost eight times (8X) the throughput (RPS) at the same 75% WAN traffic capacity.
- This 8X throughput increase was also achieved with an increase in Client LAN traffic of only three times, thus both WAN and Client LAN traffic were significantly optimized.
- Available emulated client test hardware did not permit driving the 44Mbps accelerated tests to full 75% WAN capacity; however the results show the same trends as for the 6Mbps tests, and similar percentage gains should be expected.
- The hits-per-page ratio of 3:1 (un-accelerated) was reduced to less than 1:1. This shows significant optimization in HTTP protocol, and the effects of client-side local caching of pages, content and files. The test site Home page, for example, was always read from the BIG-IP LTM with WebAccelerator cache as its content did not change.
- Average page (response) times improved dramatically, even at full accelerated load.
- Individual function response times showed marked improvements, even at full accelerated load. Search decreased by a factor of 4 (minimal page content change), Folder list open by a factor of 2 (moderate list/page content change), and document open by a factor of up to 20 (majority of documents cached by BIG-IP LTM with WebAccelerator).
- The effects of different WAN latency (90ms vs. 200ms vs. 300ms) were largely obviated, especially observing those functions that included higher page/data/file transfers.
- Although not cited in the results, a reduction in WFE CPU utilization was also noted when BIG-IP LTM with WebAccelerator was employed. This is partly due to protocol efficiencies provided by BIG-IP LTM, and in part due to client-side BIG-IP WebAccelerator caching (reducing the need for the client browser to go to the WFE for pages/files).
- The positive effects of the various technologies BIG-IP LTM with WebAccelerator employed (TCP/IP optimization, HTTP optimization, compression, caching, etc.) can be clearly seen across the data.

Internet – asymmetric BIG-IP LTM with WebAccelerator solution

While not able to provide all the capabilities and benefits of a symmetric solution, results show that deploying a single BIG-IP LTM with WebAccelerator appliance in this scenario will have measurable benefits in terms of increased throughput and in providing users with an improved experience.

- Typical throughput improvements for the 6Mbps and 1536Kbps tests approached a factor of 2 (double the throughput). Again, it was not possible to emulate sufficient load for the 44Mbps tests to drive the accelerated WAN to capacity, but the trends are similar to the 6Mbps tests and the same degree of improvement should be expected.
- The hits-per-page ratio dropped from 3:1 (un-accelerated) to about 1.2:1 (accelerated) showing a good degree of protocol optimization.
- Average page (response) times showed improvements ranging from factors of between 3 and 5 (that is, some functions took one fifth of the time).
- Client LAN traffic was reduced to 75% of the un-accelerated cases.

- A good level of compression was achieved, but note that the users' browsers are used to un-compress the data and need to be set to do so.

Summary

- The two example scenarios presented are believed to be broadly applicable to Office SharePoint Server 2007 deployments. They are especially relevant to customers considering consolidating geographically distributed farms into a few centralized data centers, and are concerned that a WAN may provide reduced performance.
- The un-accelerated tests results clearly show the potential impact of both lower bandwidth (compared to a LAN) and the performance degradation that higher latency values and increased packet loss can cause.
- The comparable results achieved when deploying a BIG-IP LTM with WebAccelerator solution demonstrate that a significant performance improvement can be achieved in terms of increased throughput, reduced response times, better WFE efficiency and overall network usage optimization.
- The BIG-IP LTM with WebAccelerator appliance is straightforward to install and configure, however HP recommends soliciting assistance from HP Services (Network Systems Group) and/or from F5 Networks Inc. to determine the appropriate BIG-IP models to suit your needs, and for assistance in the detailed configuration of the appliances.
- HP is grateful for the on-site technical assistance provided by F5 Networks Inc. during the lab setup required for the testing documented in this paper.

HP Services and F5 solutions

One of HP's strengths is to build alliances with key partners that add value to the solutions we propose to our customers. Today, HP believes that F5 Networks is providing our customers with among the best value-add solutions in the Application Delivery Controller (ADC) arena. As a result, HP has developed a strong partnership with F5 and has already developed a proven world-wide ADC solutions implementation experience. HP's selection of F5 as one of its main ADC partners is strongly reinforced by the ADC Gartner Magic Quadrant, in which F5 is positioned in the far upper right (*source: Magic Quadrant for Application Delivery Products, 2007).*

F5 and HP have a cooperative Customer Support Agreement to facilitate the delivery of best-in-class technical expertise and rapid response for HP customers' support needs. F5 and HP are proud to work together to provide customers with the best possible user experience. The terms of the agreement enable HP to resell F5's complete family of ADC solutions and to have them integrated and supported with HP networking services. As a result of this agreement, HP customers can obtain complete network design, implementation, and support services for their F5 products. HP's Global Service and Support completes the picture by providing expert networking services that help ensure availability of the customer's network anywhere, anytime the business requires.

A contact URL for HP Services (Network Systems Group) is provided in the section *For more information*. Further information relating to F5 BIG-IP solutions and products can also be obtained from F5 Networks Inc. per the contact information shown in the section *For more information*.

Implementing a proof-of-concept

As a matter of best practice for all deployments, HP recommends implementing a proof-of-concept using a test environment that matches as closely as possible the planned production environment. In this way, appropriate performance and scalability characterizations can be obtained. For help with a proof-of-concept, contact an HP Services representative. For more information, speak with an HP salesperson or reseller, or contact HP directly using the contact information available at http://www.hp.com/hps/contacts/index.html

For more information

HP ActiveAnswers resources for Microsoft Office SharePoint Server 2007	http://www.hp.com/solutions/activeanswers/sharepoint
HP ProLiant servers	http://www.hp.com/go/proliant
HP BladeSystem c-Class servers	http://www.hp.com/go/blades
HP Storage technologies	http://www.hp.com/go/storage
HP Services Network Systems Group	http://www.hp.com/go/networks
F5 Networks Inc. main page	http://www.f5.com
F5 BIG-IP page	http://www.f5.com/products/big-ip/
F5 WebAccelerator page	http://www.f5.com/products/big-ip/product-modules/webaccelerator.html
F5 White Paper on Intelligent Browser Referencing	http://www.f5.com/pdf/white-papers/eliminating-http-chattiness-wp.pdf
F5 White Paper on TCP optimizations	http://www.f5.com/pdf/white-papers/tcpexpress-wp.pdf
F5 Microsoft Solutions and Contacts	http://www.f5.com/microsoft
F5 BIG-IP SharePoint 2007 Application Ready Network Deployment Guide	http://www.f5.com/pdf/deployment-guides/f5-sharepoint07-dg.pdf

To help us improve our documents, please provide feedback at www.hp.com/solutions/feedback.

4AA1-6598ENW, November 2007