THE BEST-RUN BUSINESSES RUN SAP

F5 Application Ready Network for Enterprise Service-Oriented Architecture

F5 Networks:	**Dan Wright** **Joe Jordan** **Lori Mac Vittie** **Ron Carovano** **Randy Cleveland**
SAP Labs:	**Swapan Saha** **Joerg Nalik**

A Network Services Advisory Group Project

July 2007	1.0	
December 2007	1.1	

COMMUNITY

esc@sap.com

http://esc.sap.com

Table of Contents

1 Introduction

CIOs, application owners, and application architects are facing an ever-increasing amount of new requirements to extend their services to be available everywhere around the world. Meanwhile, they still need to consolidate data centers and business applications in order to lower costs, address new compliance rules (e.g., Sarbanes-Oxley, HIPAA), and much more. Solutions need to support the ***extended enterprise*** (Figure 1), which no longer ends at the doorstep of a headquarters data center or other traditional organizational boundaries. Company employees, business partners, and customers are all end users and demand secure and responsive access to business applications, regardless of location. In addition, the need to facilitate inter- and intra-company business application integration (enabled by service-oriented architecture (SOA)) has led to the rise in application-to-application traffic. As a result, “the network” becomes the increasingly important glue in between globally distributed business application components and end users.

SAP addresses modern business needs through enterprise service-oriented architecture (enterprise SOA) solutions, which are often delivered via the F5 Application Ready Network (ARN). In this paper, SAP and F5 present how modern application and network technologies can be combined to form productive SAP business application landscapes that perform well and securely when used globally. The work is based on the results of an SAP Enterprise Services Community Network test Lab (ENL) project in which a number of leading network and test solution vendors and SAP teamed up to build a production-like SAP enterprise SOA landscape, including applications, local end user, and remote end user simulations.

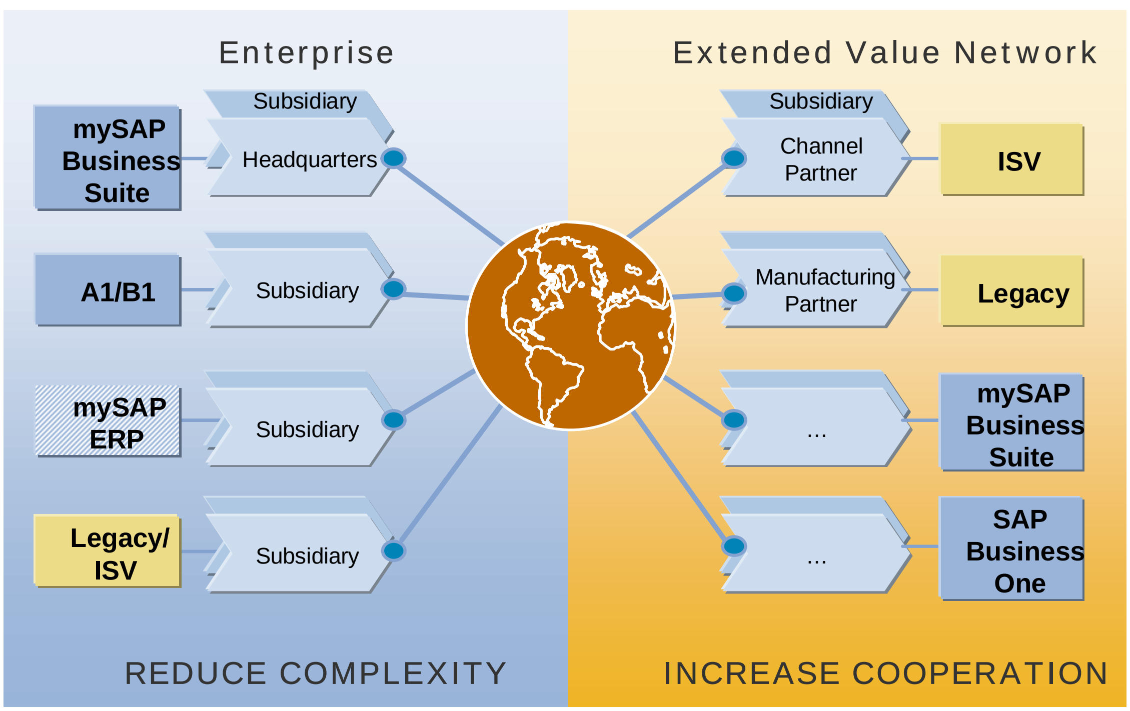

Figure 1. Extended Enterprise

When deploying SAP® applications, organizations follow SAP Best Practices recommendations to make sure these applications are installed and configured optimally. However, poor performance of intranet or Internet wide area networks (WANs) for delivering SAP business solutions to end users can be problematic, as it impedes adoption rates and introduces unnecessary delay into business processes. End users faced with a new application are already resistant to change. If the performance of the application delivery through the network is less than ideal, adoption rates and attitudes can plummet, both of which negatively impact business productivity. These types of performance issues not only impede user productivity (such as introducing unnecessary delay while interacting with the application), but can also frustrate and potentially drive away customers who interact, either directly or indirectly, with the application over consumer-grade Internet connections.

Network delay not only adversely affects the end-user experience; it can also cause problems for automated business-to-business (B2B) transaction traffic. This type of transaction is becoming more common in an enterprise SOA environment, where business processes are streamlined through the orchestration of application services that often execute without manual intervention. The efficiency gained by implementing this automation can be lost if unnecessary delay is introduced into the transaction by slower-than-expected network communication between the application services that comprise the business process.

For example, without proper architectural planning, both asynchronous and synchronous traffic can be adversely affected. Asynchronous traffic occurs when system A sends a message to system B without expecting a response. Response time improvements in such a scenario are not as relevant because system A receives a technical acknowledgement immediately and is then no longer involved in the transaction. The benefits in such a scenario are seen as a **savings in bandwidth**, which ensures that bandwidth is available for other traffic that may be more sensitive to delay. Also beneficial is the ability to enhance the reliability of messages that are forwarded through enterprise services.

Synchronous traffic, in which responses are required after processing, benefit mostly from **improved response times**. For example, a composite application might perform some SAP business logic processing and then need to do a credit check with an external agency. A delay in the response time from the external agency will result in an increased end-user response time for the composite application. The end-user response time is important not only to maintain customer satisfaction, also to increase end-user productivity by minimizing wait time.

F5's ARN is designed to optimize the enterprise SOA environment and provide cost-savings and productivity benefits for both asynchronous and synchronous traffic in three ways.

1. F5 improves the response time in synchronous application-to-application (A2A) Web-services calls, offering LAN-like speeds over WAN connections.
2. The ARN increases the reliability and availability of both synchronous and asynchronous traffic through bandwidth management and highly scalable architectures. These benefits come from F5's ability to off-load costly functions (like security) and the optimization of network resources (like bandwidth). Off-loading compute-intensive security functions, such as SSL termination, reduces the burden on application servers and allows the server to spend more resources on executing application logic. Additional benefits can be achieved by using iRules, F5's application-scripting language, to provide custom security functionality that addresses emerging threats.
3. F5's ARN consolidates multiple features into fewer appliances, reducing capital expenditures and lowering long term maintenance and associated infrastructure costs.

2 Business Use Cases

Tests were designed to 1) verify functionality and 2) provide the opportunity to feature the performance benefits of application delivery network technologies. Five business use cases were identified for potential testing.

2.1 Use Case 1: Large Enterprise SOA Objects

One of the primary-use cases was a single HTTP request invoking the transmission of a large amount of data. This scenario is played out in an enterprise SOA environment via the use of standard Web-services calls as part of an A2A or electronic B2B communication. The actual data transmitted over the network is HTTP/SOAP/XML and may contain hundreds of line items related to a purchase or sales order. In the ENL environment, it was only possible to test functional correctness with small Web-services calls. A large document-download case (wherein a 5-megabyte file is downloaded) provided a reasonable facsimile of a large Web-services call. Both large document downloads and large Web-services calls have the same characteristic of being one HTTP call with a large response from a network point of view.

2.2 Use Case 2: Application Sever Off-loading

Another area in which value can be added to SAP application installations is application server off-loading. Application servers are a core component of the application infrastructure, so it makes sense to have them focused on the most critical business logic processing tasks, especially in an enterprise SOA scenario. One of the key drivers of an enterprise SOA deployment is the need to give more remote end users access to applications for increasing business operational efficiency. The increase in users means an increase in TCP/IP connections to and in between application servers upon which services are deployed. This increase can be addressed on the application server side by simply adding more hardware servers with more application instances, to scale up processing capacity. Alternatively (or in addition to) application server up-scaling, certain processing tasks can be delegated to network appliances, which can provide even better performance at relatively low costs. Examples of such non business logic processing off-loading instances are SSL encryption/decryption, compression, static object caching, and TCP/IP connection handling. This shifts the burden of processing from the application server to specialized hardware, thus increasing the efficiency of the application server while improving the overall responsiveness of applications and services hosted on the server.

Overall, more users can be serviced by the same infrastructure, thus lowering the total cost of ownership (TCO) for the business. While all of the tests featured some level of application server off-load, a 500-user test best illustrated the benefits of this technology. This test case involved 500 individual simulated users logging onto the SAP NetWeaver® Portal component, accessing a content-rich welcome page, and then logging off after 100 seconds of "think time." In contrast to the previous one HTTP roundtrip example for document downloads and large Web-services calls, "content rich" typically translated to downloading static and dynamic welcome page data through many HTTP round-trips, which presents a different challenge to the network. Due to the occurrence of static content portions (e.g., pictures, frame elements) and dynamic content (e.g., up-to-date business data), two sub-use cases can be investigated:

1. An end user visits a Web site, downloading the welcome page for the very first time. In this case, both static and dynamic content is downloaded from the application. The static content is typically cached in the end user's browser cache.
2. During a repeat visit to the same page, only dynamic content is requested through the network from the application. The static content is taken from the prefilled browser cache from a previous visit.

The second sub-use case is much more typical for working end users that regularly access the SAP NetWeaver Portal component. However, the first sub-case leaves the "first-time impression" and is therefore also important in achieving end-user satisfaction.

2.3 Use Case 3: Optimizing User Interface and Web-Services Traffic

The third case tested general optimization of user interface (UI) and Web-services traffic that would be typical for ERP business transaction activities of end users. In an enterprise SOA application landscape, such user activities comprise:

1. Logging in to SAP NetWeaver Portal
2. Navigating to content from a composite application component
3. Retrieving and posting business data into transactional ERP back-end systems via Web-services calls

2.4 Use Case 4: Functional Correctness

The ENL landscape was built on the SAP Enterprise Central Component 2005 release running on the SAP NetWeaver platform. It was designed to simulate the manner in which SAP® applications are deployed in enterprise networks. This made it a suitable environment for conducting "functional" testing. This type of testing is meant to demonstrate that the networking technologies deployed into the environment are able to support all of the application traffic. The implemented business test scenarios contain mainstream WebDynpro UI technology as well as recently integrated partner technologies like Adobe® Forms and Adobe Flash. Furthermore, Web-services calls were implemented, as well as older SAP RFC protocols. The application servers required SAP user session stickiness and supported SSL encryption. It was critical that network vendor technology integrate smoothly into this environment.

2.5 Use Case 5: Management Tools

Deploying, managing, and maintaining network devices present unique challenges to systems and network administrators. This is especially challenging in the enterprise environment, where literally hundreds of devices can be involved, many of them being in large numbers of remote branch offices. Therefore, a set of administrative tools for network appliances are tremendously important. For example, tools may support:

- Remote datacenter and branch device configuration
- Resource consumption and event monitoring in production systems
- End-to-end troubleshooting and debugging

The ENL landscape supported the evaluation of such capabilities.

3 Test Landscape

The goal of the test landscape design was to incorporate many of the features that customers would need to use enterprise SOA applications productively. In detail this meant:

1. Using the SAP NetWeaver integration components (Figure 2) like SAP NetWeaver Portal for people integration.
2. Using SAP composite applications and enterprise SOA concept for building integrated business processes (Figure 3).
3. Building a production-like application and network infrastructure with secure access, simulated WAN links, and scalable multi-instances deployments of all SAP NetWeaver and business application components.

The first two points were fulfilled by using the SAP Discovery System software, which is a one server deployment of SAP NetWeaver, SAP xApps, and SAP Enterprise Central Component (ECC) 2005 backend applications. A number of enterprise SOA business scenarios are implemented out of the box on the SAP Discovery System which greatly reduces the effort for building a test landscape.

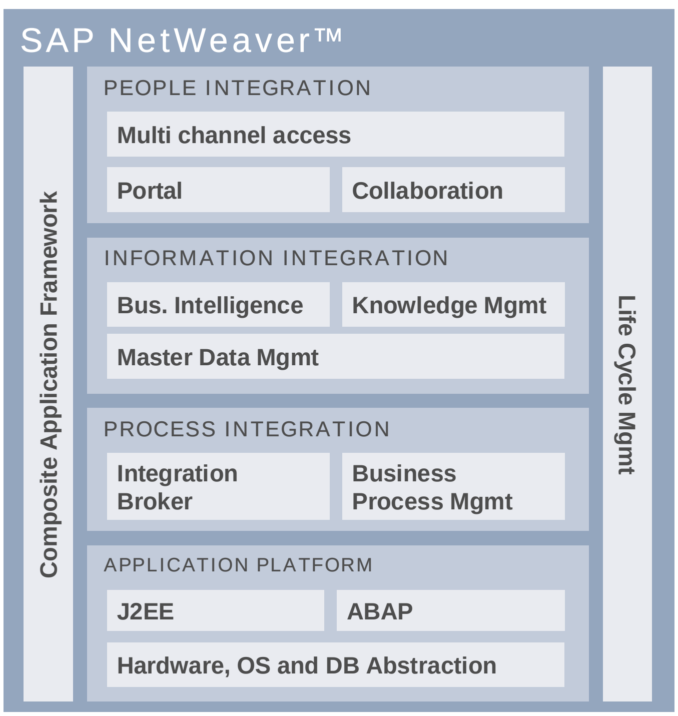

Figure 2. The SAP NetWeaver platform for running business applications.

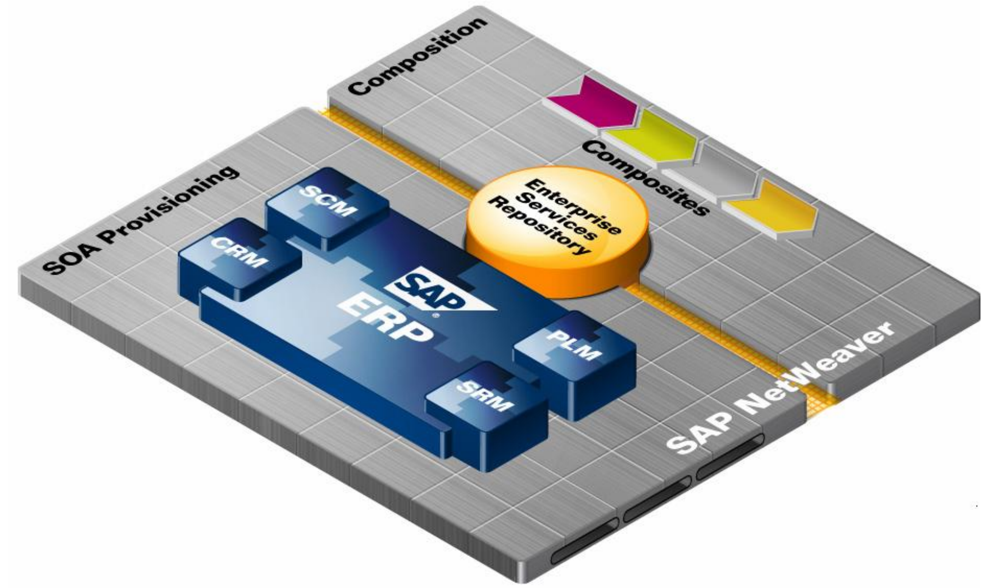

Figure 3. Building business scenarios based on SOA principles. Composite Applications are easily constructed using pre-manufactured enterprise services listed in Enterprise Services Repository to access ERP backend applications.

3.1 SAP Configuration

In order to derive the production-like infrastructure, SAP Discovery System was re-deployed onto four hardware servers as shown in Figure 4. Most components were deployed with two instances. The concept of instances demonstrated how to scale SAP landscapes to process high volumes. A real production landscape might have a very large number of instances for each component, depending on the processing volume requirements. Doubling instances for components is also a typical measure to achieve higher availability of an application landscape through added fail-over capability in case of a one-instance failure.

In order to include aspects of global use and deployment of enterprise SOA applications, the hardware servers were grouped such that they could be reside in different simulated data centers. A fifth server was added to represent the remote end user community. This crucial step added the need for a number of network services, which typically reside in an edge or De-Militarized Zone (DMZ) environment inside data centers. If end-users reside in a branch office, a small Edge/DMZ environment might also be implemented at the end-users' location.

In between datacenters and remote end users, customers would use company intranets or the Internet. In our test environment, such long distance network connections were properly simulated by using a WAN emulation appliance from Shunra© Software Ltd. For simulating end-users that execute business scenarios, which in turn trigger web service inter-component network traffic throughout the whole application landscape, HP/Mercury's LoadRunner© test tool was used.

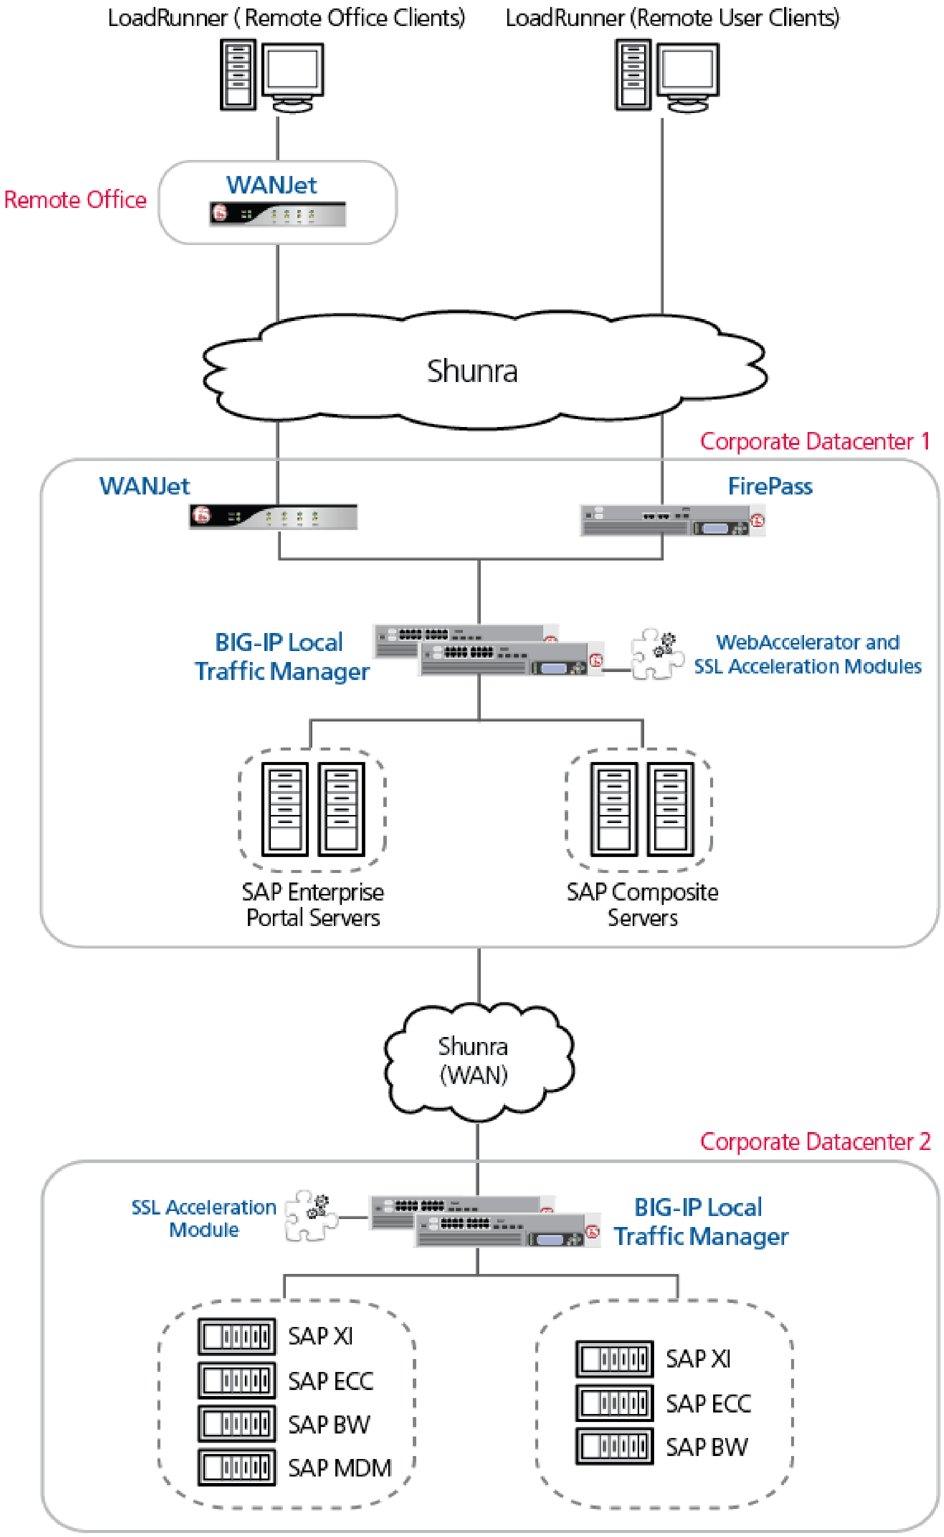

Figure 4. Test landscape (Dotted lines represent physical servers, which each contain multiple application instances.) For High Availability and for scaling processing capacity application and network components of the same kind are doubled.

3.2 F5 Configuration

F5 is an Application Delivery Network (ADN) provider. A complete ADN comprises several products that work in concert to optimize, secure, and accelerate the network for the delivery of applications and services to all locations—whether it is a desktop user in a home office, a user accessing an application via a wireless hotspot, a user in a branch office, or even another application. Just as redundant configurations in the data center provide high availability and scalability for servers, these configurations also provide similar benefits for clients. ADNs are comprised of technologies that off-load network related processing from application servers , as well as accelerate and secure the delivery of the application content to clients.

With the SAP application components in place, F5 constructed an ADN to demonstrate the additional value F5 technology brings to an SAP deployment. A redundant pair of F5's BIG-IP® Local Traffic Manager (LTM) was placed in the edge DMZ and configured to direct traffic to SAP NetWeaver Portal, Composite, the SAP NetWeaver Exchange Infrastructure (SAP NetWeaver XI) component, SAP ERP Central Component (SAP ECC) and the SAP NetWeaver Business Intelligence component services. This is the standard manner for deploying the BIG-IP LTM and provides high availability for the landscape. In addition to load balancing traffic and optimizing all TCP-based applications through F5's TCP Express technology, a number of F5's advanced features were configured and tested. These features included SSL offload and acceleration, intelligent compression, caching and WebAccelerator feature set, which further enhanced performance over the WAN by optimizing the HTTP protocol. WebAccelerator employed a number of client-focused features to improve and enhance communication between the browser and Web-based applications.

All of these advanced features, when combined with the BIG-IP LTM's protocol optimization capabilities, provide end-to-end optimization and performance enhancements of networks for Web-based applications and services, such as those deployed within an enterprise SOA infrastructure. This allows F5 to offload some of these processor-intensive tasks and free server-side resources, such that more hardware server resources become available for the SAP applications to concentrate on business logic, thus increasing the capacity and efficiency of those servers.

F5's FirePass® SSL VPN device was also installed in the edge DMZ area to facilitate testing of its secure application access technologies. The FirePass controller is an appliance built to provide secure application access from anywhere, via any device. In our testing, we deployed a single FirePass device, although multiple FirePass controllers can be deployed to enhance reliability. These devices can be located in the same data center or placed in each of the enterprise's primary datacenters. Tests were conducted in the Enterprise Services Community Network Test Lab (ENL) environment to ensure the compatibility of the FirePass controller's various application access methods with SAP applications.

The final piece of this application delivery network was a pair of F5 WANJet® appliances. These devices are built to deliver Local Area Network (LAN) like performance over the Wide Area Network (WAN). The addition of the WANJet devices provided advanced symmetric acceleration that would be applicable to connections between central offices and remote offices or business partners. In this testing scenario, one device was placed in the edge DMZ for SAP NetWeaver Portal and composite services servers. The peer device was placed in the edge DMZ of the Client's "Remote Office" (see Figure 4). These devices can also be installed in a hub and spoke deployment for enterprises with multiple remote offices. The combination of these appliances provided significant improvements for both client-side and server-side application performance. WANJet appliances also played a security role in the deployment. The communication between the two devices happened over an SSL tunnel. This means that unencrypted traffic such as HTTP directed over the WANJet link ended up being encrypted as it transverses the WAN.

Figure 5 summarizes the how SAP applications benefit from F5 solutions.

BENEFIT FOR SAP APPLICATIONS	F5 SOLUTIONS
Optimization • Offload processor intensive network related tasks that are not core business logic functions • Accelerated WAN performance user experience • Reduction in network bandwidth usage	**BIG-IP LTM with WebAccelerator Module** • Load balancing • TCP optimization • Connection multiplexing (OneConnect) • SSL offload and acceleration • Intelligent compression • Dynamic caching • Intelligent browser caching **WANJet** • Bandwidth management, compression, and Quality of Service • TCP optimization
Availability • Application availability through the network • Application level resource control	**BIG-IP LTM** • Advanced application health monitoring and reporting for load balancing optimization. • Server and system level resource management • Advanced load balancing based on multiple metrics, such as response time, number of connections, CPU and memory utilization • Layer 7 rate shaping
Security • End-point security • Secure access without special client software • Protection of corporate assets • Unified access across client devices • Denial of service protection • Network security as part of overall Application security • Secure data replication	**FirePass (SSL VPN)** • Client integrity verification and remediation • Cache cleanup and download blocking • Protected workspace **BIG-IP LTM** • Resource cloaking • Cookie encryption • Network-level protection **WANJet** • Site-to-site SSL encryption

Figure 5. Benefits Summary

4 Test Methodology

SAP and F5 drafted a number of test cases to script interaction with the SAP applications and F5 ADN devices over a set of different simulated internet links. ***The goal was to illustrate how different network connection qualities might affect the performance of network for the applications and to demonstrate the ability of the network appliances to increase availability and security.***

Therefore, tests accounted for four major independent parameter sets:

1. Use of different SAP business scenarios
2. End-user conditions
3. Network conditions
4. Application of F5 technologies

Standard measurement tools and procedures were employed to ensure that results could be accurately captured and compared.

4.1 SAP Business Scenarios

There are literally hundreds of different business scenarios with many variations which might be implemented in SAP customer's enterprise SOA landscapes. However, the WAN induced challenges to performance and security are typically very similar from scenario to scenario and fall into only a few categories. Therefore, only very few scenarios were chosen and scripted for use with HP LoadRunner. The chosen scenarios cover the common aspects of:

1. end-users logging in to SAP NetWeaver Portal as launch point for all further activities.
2. one example for typical transaction like processing in SAP systems
3. one example of use of composite applications, a central architecture element of enterprise SOA
4. large data object downloads as it occurs for document management systems but also is the characteristics of large Web Services calls

A typical scenario starts with and end-user requesting the SAP NetWeaver Portal login page, then submitting login user/password credentials and getting EP's first page which is commonly referred to as "welcome" page. From here a user navigates typically through a few steps to a more specific inter-action which then would also trigger activities in other enterprise SOA landscape components like the composite applications or backend application components.

Though each business simulation script of for HP LoadRunner contains many steps, only few key steps are pointed out in this document in order to keep it comprehensive. The effects of WAN and F5's WAN improvement technologies can be sufficiently shown and explained with those key steps. Some of them are:

- **Sales Order Creation:** A salesperson takes a customer's order and records the order. The salesperson then enters and submits the customers purchase information. This is a typical SAP transaction like activity. It demonstrates how a sales order is handled in an enterprise SOA setting where through the composite application component a web service is created and sent to the SAP Enterprise Central Component backend through the SAP NetWeaver XI message hub.
- **Get Supplier Info by Product:** This activity will demonstrate how to use cross-component Business Process Management (ccBPM) in the SAP NetWeaver XI message hub to split a message and send it to the respective receivers using a web service in par-

allel. The receivers respond with confirmations. This scenario is an example for business to business (B2B) interactions. "Get Supplier Info by Product" is the asking of multiple Suppliers for a Quote on products. In order to get the quotes from suppliers, technically a web service call is made from the composite component to SAP NetWeaver Master Data Management component, business data integration component. A list of suppliers and their contract information for a list of product(s) is returned.

- **5 MB Knowledge Management Document Download**: In contrast to so called page-downloads, which typically consist out of a number of client-application http(s) round trips, there is another category of requests which consists out of exactly one http(s) request, which delivers a large, often many MB size data object back. An important example out of the enterprise SOA world are application-to-application (A2A) web service calls, which transport the data of whole business objects, like for instance a purchase order. Very often such business documents have hundreds of line items which leads to the MBs of size when transported in XML format. Another example is just any kind of documents which could be retrieved from the Knowledge Management component in SAP NetWeaver Portal. In the ENL a 5MB Power Point from Knowledge Management was chosen as a simple test case for the one http(s) request/large data object category.

4.2 End-user Conditions

Tests were designed to simulate a single user stepping through a defined sequence of data input and click-steps in a browser. Additional tests were composed with multiple users simultaneously accessing the system for an extended period of time. The multi-user tests were built to illustrate the stability of the system, as well as to demonstrate some of the positive effects of offloading SSL and other processings from the application infrastructure.

Each test case required running the tests with the HP (former Mercury) LoadRunner® client cache being cleared between each iteration, and a second complete run without clearing the client cache between iterations. The first test run was intended to simulate a client visiting the application for the first time. The second test run was intended to represent users that have logged onto the application in the recent past, so much of the framework was already stored in their browser’s cache. As a result, only the dynamic content was downloaded, which generally involved transmitting less data and fewer requests. When performing automated tests with the LoadRunner tool, efforts were made to run a few of the different scripts at the same time. This allowed for completion of more tests in a shorter period of time

The long term multi-user tests involved a ramp-up period of about four hours. Over that time, the number of virtual users was gradually increased, and the client caches were populated for each of the users. During the slow ramp-up the gradual application server and network resource consumption increase could be monitored Once 500 virtual users were initiated, the test ran for another one hour without clearing the client browser cache. We chose to ramp to 500 users, using a 100 second Think Time users take in between performing requests, based on the baseline performance of the ENL SAP environment without any network appliances in the data path. Under these conditions the SAP systems alone run under high load but were not yet overloaded.

4.3 Network Conditions

Four different network test links were evaluated (Figure 6). These links were enforced via a Shunra WAM emulating appliance. Before running each test, the Shunra appliance settings were updated to be consistent with one of the four different test links. The smallest bandwidth was meant to represent a home user in North America accessing a remote application via DSL. We also tested with a link of 45 Mbps and 60ms latency (30ms each way),, representing a connection between two major corporate offices on the same continent. Two additional long haul bandwidth links were tested, representing a connection from an office in North America to an office in

Europe or Asia-Pacific. The long haul measurement with 1% Packet Loss took into account that public network quality to some countries is not optimal. Ping times were measured between configuration changes to verify proper latency simulation.

Link	Bandwidth	Latency (RTT)	Packet Loss
Link 1	**786 Kbps**	**60 ms**	**0%**
Link 2	**45 Mbps**	**60 ms**	**0%**
Link 3 (long haul)	**45 Mbps**	**300 ms**	**0%**
Link 4 (long haul)	**45 Mbps**	**300 ms**	**1%**

Figure 6. Network test links (RTT=Round Trip Times)

4.4 Application of F5 Technologies

Several different configurations of network equipment were tested. There was a baseline case where SAP handled the load balancing for SAP NetWeaver Portal, composite services, gzip compression and SSL handling and encryption were done by the SAP application servers, without any further optimization from any network appliances. Then we tested the application infrastructure by adding each of F5's acceleration products one by one, to understand the contribution of each product to overall system performance. Host file changes were used to guide the HP LoadRunner client directly to the SAP servers or through the F5 appliance infrastructure. This provided a simple way to compare baseline performance versus the performance of the applications when delivered using F5 technology.

4.5 Measurement

When all of the necessary configuration changes were put into place to support a given test case, the LoadRunner script was initiated. The data was collected from the LoadRunner statistics and analyzed after the test run. Other tools, such as network sniffers, were also employed when setting up the testing environment to verify the accuracy of the results. User Think time was subtracted in the course of processing the data, so only the time between user requests and receiving the complete response was measured.

5 Test Results

Practical constraints (such as time, limitations of the test landscape) did not permit the opportunity to comprehensively test every combination of end-user conditions, network conditions, and application of F5 technologies. Nonetheless, significant benefits were measured, which support general observations about network enhancement for SAP applications when delivered over the WAN. Qualitative and quantitative highlights are summarized below.

5.1 Qualitative Results

Wide Area User Performance Improvement

SAP applications were shown to perform extremely well over the LAN. In this environment, latency was practical zero and available bandwidth was high (100Mbps). However, when testing enterprise users attempting to access the applications over the WAN, performance began to degrade rapidly. Sometimes this degradation was due to constrained bandwidth between the communicating parties. The degradation may also have been related to the latency imparted on the system as a result of the (simulated) geographic distance between the clients and the servers. Reducing the number of round trips required to perform a given transaction with an application is paramount to achieve acceptable application performance for geographically dispersed WAN users. F5's products provide a number of ways to reduce the number of round trips and overcome the impact of conducting business transactions over the WAN.

More efficient communication between remote offices

F5's WANJet® is an appliance-based solution that delivers LAN-like application performance over the WAN. Because of the impact of TCP Window Size and the limits it places on the amount of data on the wire, adding bandwidth is often ineffective at reducing application interaction time. TCP throughput degrades significantly on the WAN, particularly on high-latency, intercontinental links. To overcome these inherent protocol limitations, the WANJet device employs adaptive TCP optimization (which combines session-level application awareness, persistent tunnels, selective acknowledgements, error correction, and optimized TCP windows) to fully utilize available bandwidth. This enables WANJet to adapt, in real time, to the latency, packet loss, and congestion characteristics of WAN links, and accelerate virtually all application traffic. These dramatic performance increases can be seen in Figure 13.

F5's technology also allows SAP's application servers to focus on high-value business functionality instead of spending valuable processing cycles on things like caching content, compression, SSL en/decryption, TCP/IP connection handling, etc. All of these capabilities are vital for productive enterprise SOA landscapes, but are provided with a much lower TCO through F5 technology. For example, the testing showed that by allowing F5's BIG-IP LTM to optimize TCP connections, 1000 client connections were reduced to 50 server connections. This increased the efficiency of the operating system and lowered the CPU utilization of the application servers, allowing the applications to serve more concurrent users with better response times.

For end-users visiting an application web page for the first time (with empty browser caches) F5's BIG-IP LTM was also able to offload cacheable content, reducing the number of web objects requested from the application servers. When any web object was requested from the application servers, F5's TCP Express feature-set enabled the SAP NetWeaver Application Server component to rapidly hand the requested object to F5's BIG-IP LTM device and then move on to other requests. F5's BIG-IP LTM then managed the delivery of this content to the client over the slower WAN. Additional SAP NetWeaver Application Server cycles were then available to service other requests once they were no longer burdened with managing traffic to slow clients.

Reasons why WAN response times decay:

Latency times: The TCP/IP protocol chops a large data stream into packets (L4-OSI function). After each packet sent, the receiver needs to send back an acknowledgement to the sender before the next packet can be sent. Such Round Trip (RT) can't be faster than the signal travel speed on physical lines (L1-OSI) which is somewhat less than the speed of light. This leads to a Round Trip Time (RTT) delay. A small packet network ping can be used to measure RTT. The packet size is called "TCP/IP window size" and is typically 64KB for Windows-OS. **Example: a 5MB document needs about 80 RT to be delivered, leading to an added response time of about 80*<ping time>.**

Bandwidth times: The finite bandwidth of transmission lines adds response times too. **In the DSL (768Kbps) case this can add about 50sec response time for a 5MB document.**

Packet losses: Lower quality network WAN lines exist in some parts of the world and can cause occasional TCP/IP packet losses. While loss rates are often low, like 1%, they have a drastic response time effect. The TCP/IP window/packet size is dynamic and is reduced as a precaution when packet losses happen. **This has the effect to increase the numbers of round-trips and therefore latency time losses, often by big factors.**

Figure 7. WAN response issues

Additionally, the job of handling the complex mathematical computations for SSL was assigned to the dedicated hardware present in the F5 BIG-IP LTM device, further reducing the workload of the application servers. Studies indicate that SSL processing consumes approximately 30% of a server's resources. Thus, assigning SSL processing to an external device can ostensibly reduce the number of servers necessary to service the same number of users from three to two. This reduction lowers TCO by reducing the hardware, software, licensing, and maintenance costs of additional servers and application instances. It also yields a reduction in the maintenance and costs of the digital certificates required to support SSL by centralizing certificate management. Offloading SSL (from individual application servers to F5 BIG-IP LTM) decreases the time and cost to deploy servers by removing the complex task of installing and configuring SSL support and certificates on each individual server. It also removes the need for application-specific configuration or specialized code to manage secured connections between B2B focused services in an enterprise SOA environment. F5's dedicated SSL hardware processing components also offer businesses the freedom to implement SSL without paying a performance penalty. The results of these F5 optimizations can be seen in Figure 11, where F5 reduced SAP CPU utilization from 68% to 38%.

5.2 Quantitative Results

768 Kbps Link Highlighted Results

The 768 Kbps test case was created to demonstrate application performance for clients accessing the SAP application from a remote office with very low available bandwidth or a home user. The latency was set at 60ms round trip time, much like the latency between two major cities. While the latency was not as extreme as in some of the other test cases, this case had the lowest bandwidth, therefore significantly impairing the performance of the application. Downloading the application framework components, along with the documents, took significantly longer than the same transaction on the LAN. The extra time represents a decrease of efficiency and an added cost for conducting business over this link.

F5 demonstrated significant application performance improvements with the 768 Kbps link (Figure 8). Although the base response time (blue bar—representing SAP applications on the WAN) for the pages in question was very fast, acceleration technologies employed by the WebAccelerator further enhanced the end user experience due to its unique browser-focused performance enhancement feature set. In these instances, performance gains (red bar—representing SAP applications with F5 on the WAN, and the yellow bar— representing SAP applications with F5 on the LAN) made by F5 were primarily from a combination of rapid page generation due to patented Application Smart Caching (ASC) and round trip latency mitigation from Express Connect technologies. The yellow bar shows performance of SAP applications on the LAN (without F5 technology) and thus represents a theoretical performance limit for SAP application.

Briefly, ASC understands an application's high-level logic (what can and cannot be cached, what events cause invalidation, etc.), thus eliminating repeated processing of complex web requests. ASC enables WebAccelerator to decide when to invalidate objects and how to identify reusable pieces of content, thus accelerating performance. Likewise, Express Connect upgrades the standard browser-to-server connection by handling multiple requests and responses in parallel, often without involving the origin server, while optimizing server scalability and bandwidth capacity.

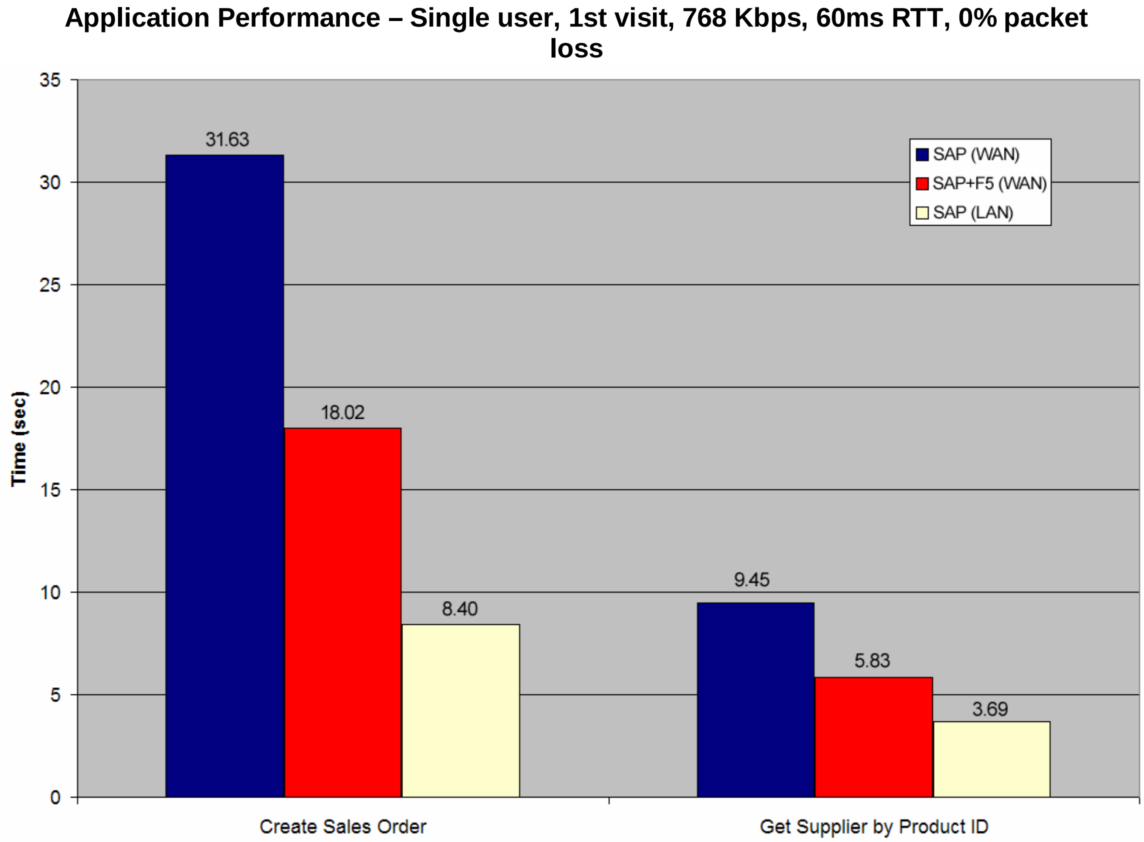

Figure 8. Application Performance improvements with F5 and a 768 Kbps link

Performance gains for document download on the low speed 768 Kbps network were substantial (Figure 9). These gains came primarily from the Express Compress advanced compression capabilities and TCP optimizations in BIG-IP LTM, WebAccelerator and WANJet. Express Compress utilizes gzip technology, but instead of caching the compressed object, WebAccelerator breaks up the content into a compressed instruction set that can be used to re-generate the content at near-line speeds. This allows WebAccelerator to modify the compressed content without having to uncompress and re-compress the entire object.

For both tests shown in Figure 9, an additional F5 configuration was tested in which WANJet was configured to handle SSL encryption (green bar) instead of BIG-IP LTM (red bar). Using the WANJet and its Transparent Data Reduction technology, further accelerated performance to nearly the theoretical performance limit (yellow bar) was found. Transparent Data Reduction utilizes a two stage compression process to maximize bandwidth savings while minimizing processing latency. The first step of the process examines the transmitted data to determine if any part of it has been previously sent. If so, the previously transmitted regions are replaced with references. The second step further compresses the data through the use of dictionary based compression and advanced encoding schemes.

Application Performance – Single user, 1st visit, 768 Kbps, 60ms rtt, 0% packet loss

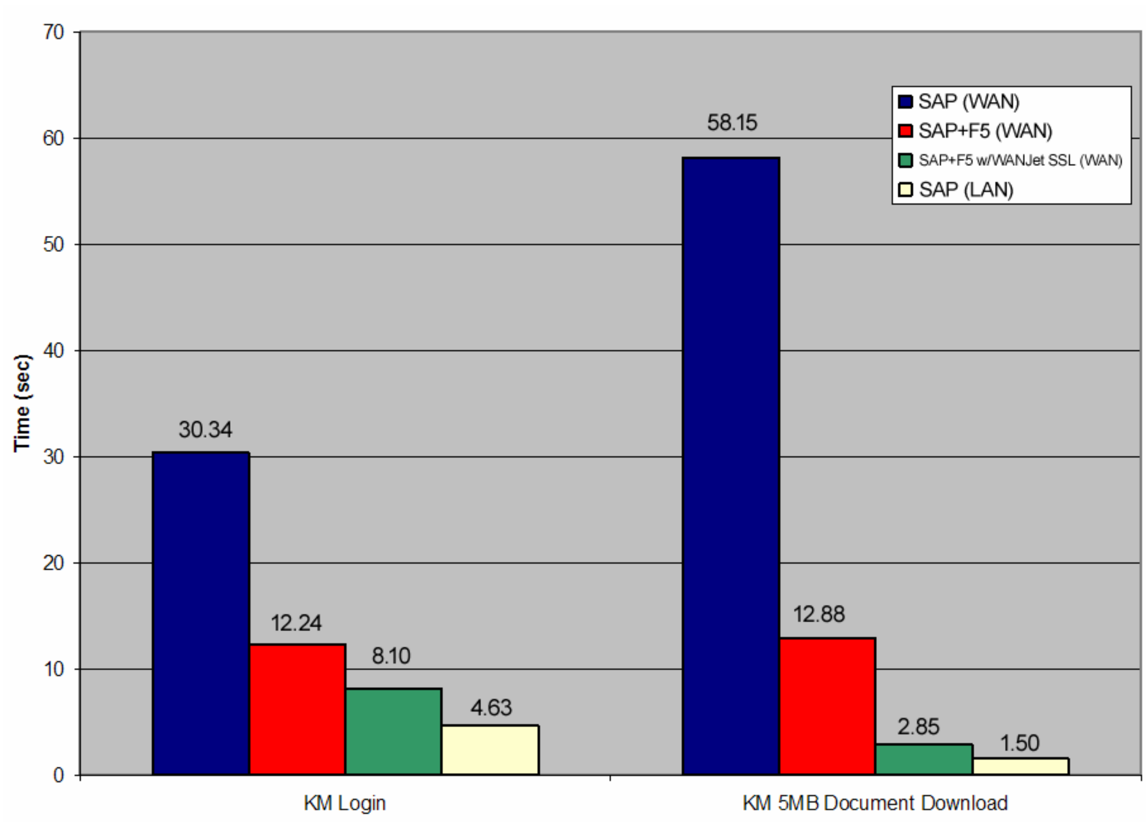

Figure 9. Login and Document Download Performance improvements with F5 and a 768 Kbps link

45 Mbps Link Highlighted Results

The Multi-User test involved virtual users logging on to SAP NetWeaver Portal with SSL, receiving their homepage, and then logging off without clearing their browser cache. The virtual users would then repeat the process until the test completed. The test case involved clients accessing the SAP application over a 45 Mbps link with 300ms Round Trip Time and no packet loss. The 45 Mbps test case is, for instance, representative of a high bandwidth connection between a North American office and a manufacturing plant in Asia-Pacific. A document download that may have only taken seconds on the local area network, may now take several minutes to complete. Making improvements in this area can have significant positive effects for the application clients.

The test was run for one hour at sustained maximum load to demonstrate stability of the system. Figure 10 shows the average response time for several transactions. These were repeat transactions and are
representative of many of the transactions that occur in the daily utilization of an application.

Application Performance – 500 Users, repeat visit, 45 Mbps, 300ms RTT, 0% packet loss

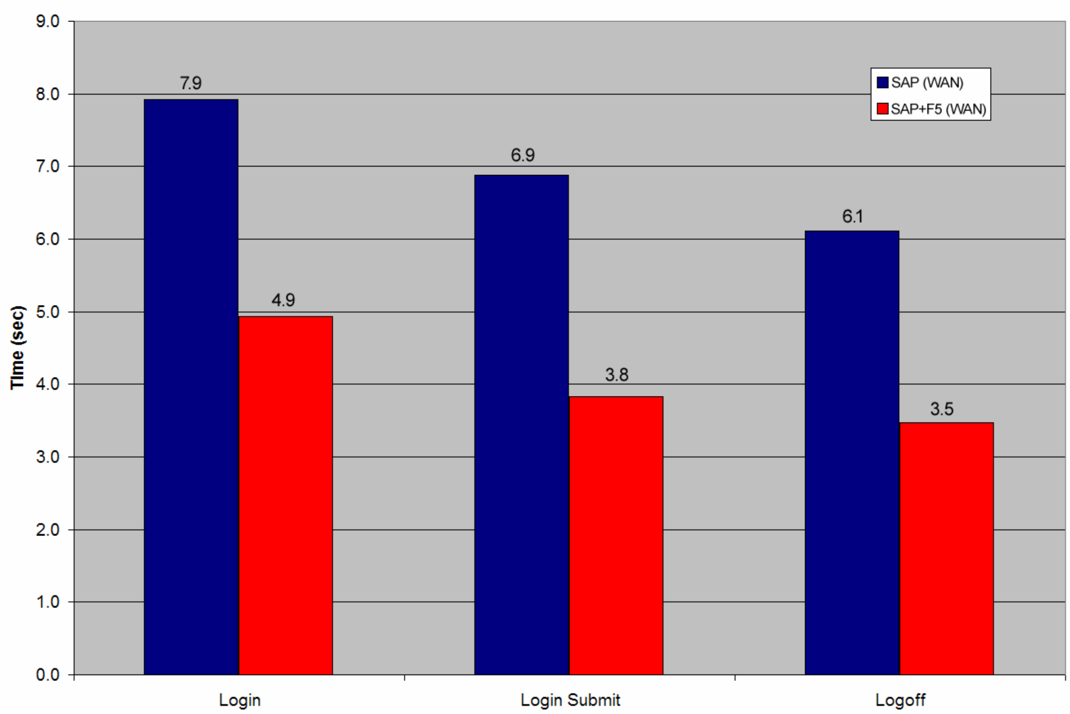

Figure 10. Page load time improvements with F5 for 500 users

Figure 11 shows how the average CPU utilization of the SAP NetWeaver Portal application servers was reduced with the implementation of F5 technologies. 500 users were ramped up over a four hour period, with an additional hour of steady-state interaction. The average was taken after each of the 500 users had loaded content into their browser caches. As a result, all of the requests from those 500 users were requests for dynamic data. The requests were secured with SSL, which was decrypted by the F5 appliance. The F5 appliance also aggregated the requests from 1000 TCP connections down to approximately 50 server side connections. A number of F5 features working together enabled this significant reduction in the load on the servers.

CPU utilization as load steadily increased to 500 users

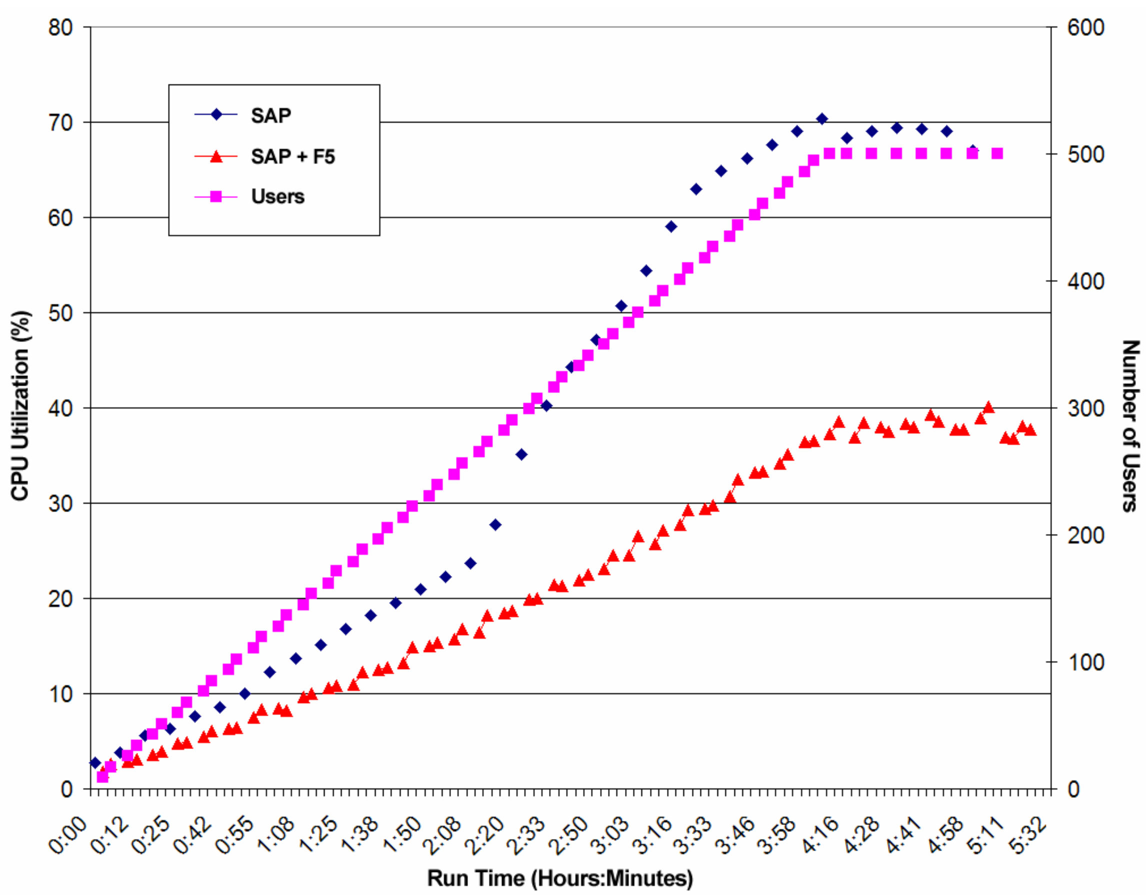

Figure 11. SAP NetWeaver Portal CPU usage as number of users increases

The following test results are for the same 45 Mbps link, but with the addition of some packet loss. This test helps illustrate the impact of packet loss on the transfer of data. Similar to the other tests, the F5 products showed performance gains across all user transactions. Unlike the 768 kbps tests though, first visit performance improvements were primarily due to TCP Optimizations to overcome high latency and packet loss. This was especially evident for the document download case where a larger number of packets are lost over the duration of the transaction.

Application Performance – Single user, 1st visit, 45 Mbps, 300 ms RTT, 1% packet loss

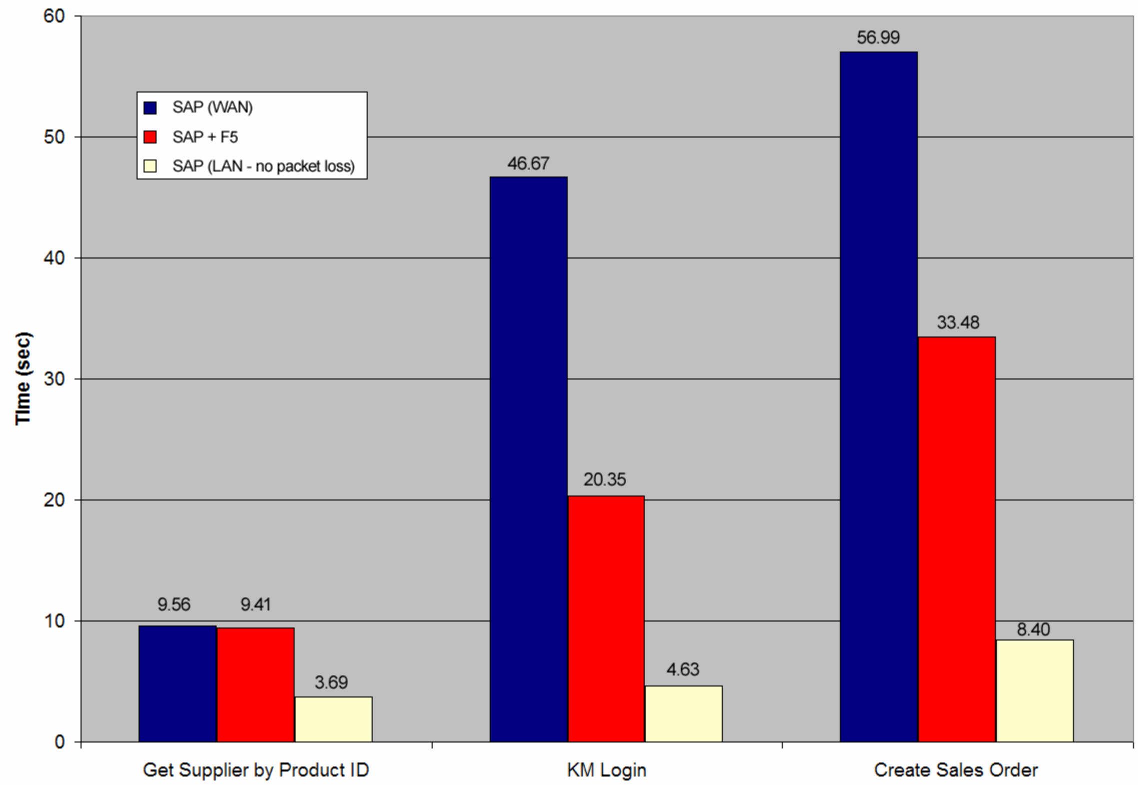

Figure 12. Application Performance improvement with F5 and a 45Mbps link

Document Download Performance – Single user, 1st visit,
45 Mbps, 300 ms RTT, 1% packet loss

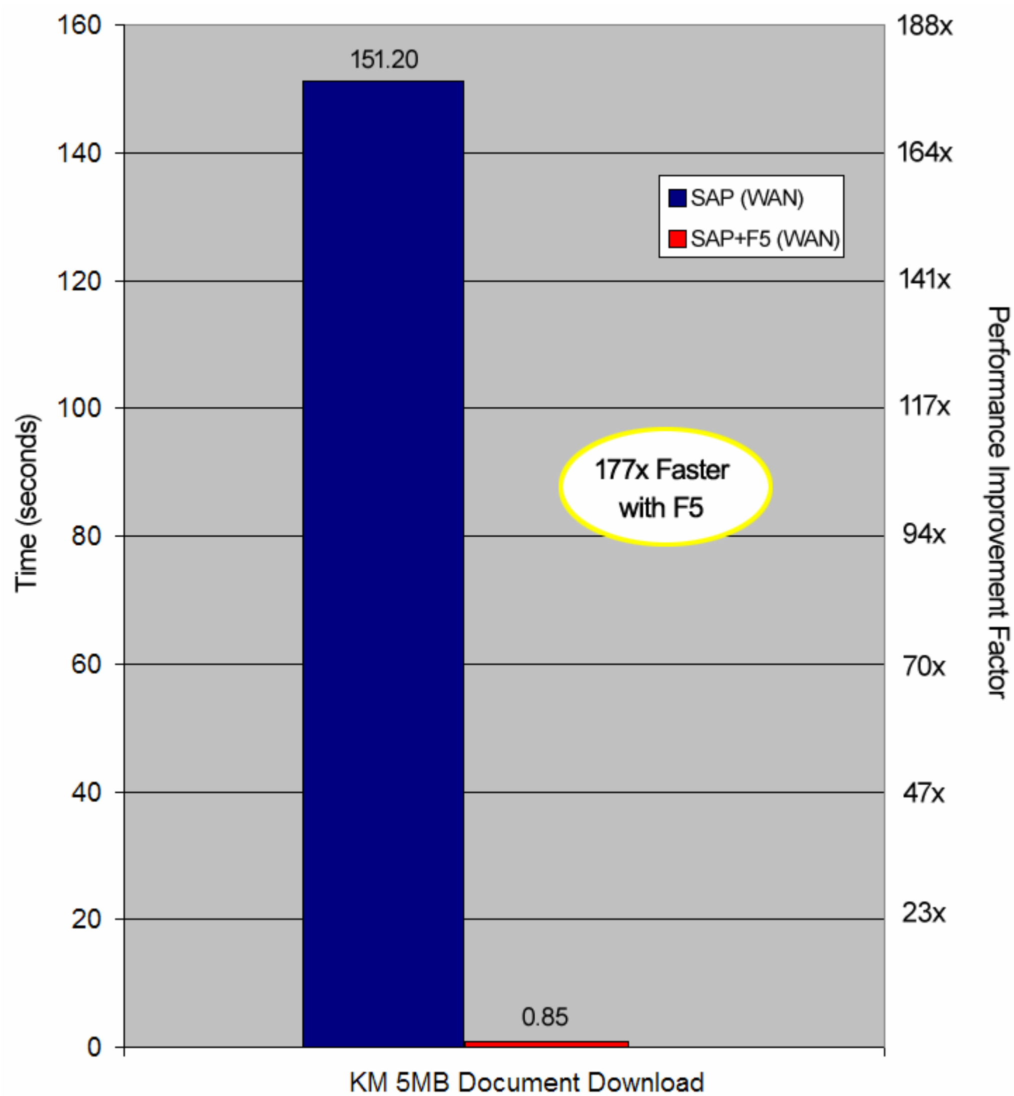

Figure 13. Document Performance improvement with F5 and a 45Mbps link

**The extremely short time required to download this single document was due to the advanced TCP optimizations in F5's WANJet device and was achieved even without also leveraging F5's Transparent Data Reduction technology.

6 Unquantified F5 Benefits

6.1 Browser Cache Optimization

The default configuration of SAP NetWeaver Portal makes significant use of HTTP cache control headers to optimize repeat visit performance. Static objects, such as images and JavaScript, receive a configurable expiration time (SAP recommends one week) so that they are not downloaded time and again during the expiration period but served to end users out of their browser cache. However, in order to be able to rollout changes of static content to end-users the expiration time can't be set too long. With WebAccelerator's Dynamic Content Control technology, this duration is safely extended out months longer, with no impact on data accuracy. How does this work? Briefly, WebAccelerator makes requests to the origin server and adjusts tags on pages sent to the client. Thus, when new content is available, tags change and the client automatically requests new content from the origin server, rather than retrieving it from the local cache. The potential performance gains created by this were not demonstrated in the context of the ENL engagement, but have been vetted with actual customer deployments.

6.2 Universal Application Access

F5's FirePass SSL VPN appliance enables secure remote access to an enterprise's applications from anywhere. FirePass also provides services to guarantee that devices logging onto the enterprise's network meet the organization's security requirements, for examples, having a working and current anti-virus protection in place. If the devices do not meet the requirements, they can have access refused or perhaps be given access to only a limited number of resources. The motivation for implementing this technology is to protect application servers from viruses and malware that can cause unexpected down time or consumption of resources all of which impede the performance of the application for all users. FirePass can even use existing enterprise authentication services, thus streamlining its installation into the enterprise network. The FirePass device's client also provides for TCP compression and additional caching to enhance performance for the remote users accessing the enterprise network.

FirePass Application Tunnels are designed to provide controlled access to an application. Administrators can configure secure, application-level TCP/IP connections from the client to a specific set of IP addresses and ports on the network. The FirePass client is still used, but the user can only create connections to the specific resources on the enterprise's network. Implementing Application Tunnels for business partners make a lot of sense, because it allows an enterprise to expose access for a specific application to a small group of users without granting those users direct access to the enterprise's entire network.

All of this functionality was tested in the context of the ENL engagement to ensure F5's compatibility with SAP applications.

6.3 Management

All of F5's products can be managed via a web based interface. Security for this interface is provided via SSL. Device configuration and statistics are available through this interface. The majority of F5's products can also be administered via a command line interface. This interface uses SSH for security. Some platforms also have "lights out management", providing the ability to remotely power on or power off hardware, reboot, and perform routine administration.

F5's standards-based iControl API provides a programmatic interface for managing the configuration of F5's products. This API is based on WSDL and SOAP and integrates smoothly into new or existing SOA environments. iControl can be utilized by application owners to make changes to the network's behavior as a response to additional demand or during maintenance windows. iControl applications can be constructed in a variety of programming languages, giving administrators the choice to choose the language with which they are most comfortable. Example iControl applications can be found on http://devcentral.f5.com.

F5 also supports virtual administration domains. This allows for a single device to be managed by multiple application teams without interference. Every user can be assigned to specific Administrative Domains that define which objects are visible to a specific user. Multiple levels of access are also definable for each user, ranging from basic Read Only users who can log on to the devices to monitor status of Virtual Servers and traffic quantities, to full Administrative users capable of making configuration changes to every object on the device. This reduces the time spent in meetings, tracking down appropriate administrative personnel, and improves the ability of application administrations to manage applications as needed. This streamlines the business process and improves the efficiency of operational personnel.

For those organizations with larger numbers of appliances, F5 offers Enterprise Manager. Enterprise Manager is an appliance that gives administrators the power to centrally discover and maintain all F5 devices in their network. Enterprise Manager enables the archival and safeguarding of device configurations. Enterprise Manager is also able to configure new devices and roll out software upgrades and security patches from a central location without manually working on each device. This reduces labor costs for managing multiple F5 devices. The figure below shows an example of user interface for Enterprise Manager.

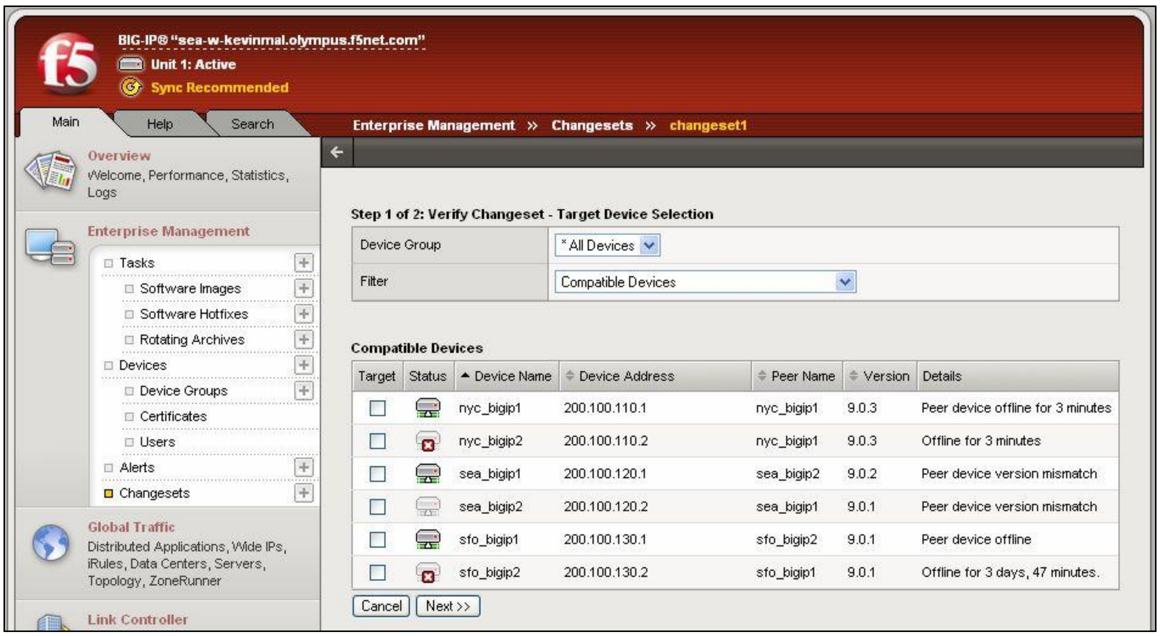

Figure 14. Enterprise Manager User Interface

6.4 Application specific guidance

F5's Solution Center is a resource providing application specific configuration guidance for various applications including SAP. These materials provide the details on what configuration options to employ for best performance. These documents are available in the F5 Solution Center at http://www.f5.com/solution-center/partner-showcase/sap.html.

6.5 Capabilities not featured in the ENL environment

The WANJet device's QoS feature lets IT managers tune their global networks for round-the-clock application performance. IT resources are logically grouped and managed according to business rules. As a result, bandwidth can be preserved for real-time mission critical SAP applications. This translates into a predictable user experience for application performance, regardless of WAN conditions, time of day, or geographical location.

Identical applications installed in different data centers can be intelligently load balanced via F5's BIG-IP Global Traffic Manager (GTM) module. This product uses DNS technology to direct users to a specific data center based on availability and performance. Administrators can also direct requests by proximity to deliver clients to the closest data center.

F5's BIG-IP Link Controller module is designed to provide a more flexible implementation for load balancing outbound requests from data centers. Link Controller frees the enterprise from the complexities of BGP peering between competing carriers. With F5 Link Controller technology, an organization has the choice of aggregating multiple small connections together rather than having to invest in a single high bandwidth connection. This frees businesses to expand their service as they grow. The Link Controller seamlessly monitors availability and performance of multiple WAN ISP connections to intelligently manage bi-directional traffic flows to a site, providing fault tolerant and optimized Internet access. F5 devices detect errors across an entire link to provide end-to-end, reliable WAN connectivity. In the event of a failure, traffic is dynamically directed across other available links so users stay connected.

The BIG-IP Application Security Module (ASM) runs on the BIG-IP application traffic management platform, providing robust application security with BIG-IP traffic management capabilities in a single system without the need to buy or install more hardware. The ASM provides application layer protection from both targeted and generalized application attacks to ensure that applications are always available and performing optimally. Together, BIG-IP with ASM offers a complete, robust solution that reduces box clutter, lowers maintenance and management costs, and provides a new level of proactive application protection while ensuring exceptional application performance.

iRules are one of the most unique capabilities of F5's products, which can greatly enhance application and network performance and security. iRules are a feature of the BIG-IP platform that allows developers and network professionals to create and customize policies that can enhance application flow and availability. F5's DevCentral serves as an online community for building and optimizing iRules. One example is cookie encryption, which provides the ability to secure and protect cookies from outsiders (http://devcentral.f5.com/wiki/default.aspx/iRules/EncryptingCookie.html).

7 Summary

Businesses want to consolidate their IT operations to reduce costs and at the same time allow a growing number of globally distributed end-users access to business applications. Still, businesses often need regional datacenters, want to outsource operations to hosting providers, and seek to integrate remote applications from business service providers (e.g., credit check service) and business partners. Consequently enterprise SOA applications are being deployed in a more geographically dispersed manner as businesses become more globally connected through the underlying Intra- and Internet infrastructure. Disparate requirements of IT consolidation versus globalization of the business often compromise the performance of business applications whose rapid execution is critical to maintaining a competitive organization. F5 and SAP set out to construct a model application infrastructure, and use this as a platform for investigating ways of counteracting the common performance shortcomings of applications deployed over the Internet, while ensuring the much needed security and reliability of business communications.

The benefits of F5's application acceleration technology were verified with the tests that quantified the performance impact of the various products for clients accessing SAP applications over the WAN. F5 Application Delivery Networking devices provided performance approaching LAN speed for users accessing the applications over low bandwidth and high latency links. Security, faster end-user response times, and cost savings through offloading network related tasks from application servers to F5 network devices were demonstrated, as well.

Together, F5 and SAP have thoroughly tested and demonstrated the benefits of the F5 Application Ready Network. Optimum performance of applications and the network leads to increased user adoption, satisfaction, and productivity. Ultimately, this means decreased costs and increased ROI for CIO's and Business Application owners.

8 References

SAP References:

SAP Developer Network (SDN)
http://sdn.sap.com/

SAP Enterprise Services Community (ESC)
http://esc.sap.com/

SAP Discovery Server
https://www.sdn.sap.com/irj/sdn/developerareas/esa/esadiscovery

First ESC Network Group paper
https://www.sdn.sap.com/irj/servlet/prt/portal/prtroot/docs/library/uuid/805d8c2d-0e01-0010-a694-a94109e88f2a

F5 References:

BIG-IP Local Traffic Manager (LTM).
http://www.f5.com/pdf/products/big-ip-8800-ds.pdf

FirePass SSL VPN.
http://www.f5.com/pdf/products/firepass-overview-ds.pdf

WANJet.
http://www.f5.com/pdf/products/wanjet-ds.pdf

Dynamic Caching
http://www.f5.com/pdf/white-papers/dynamic-caching-wp.pdf

Transparent Data Reduction.
http://www.f5.com/pdf/white-papers/wanjet-tdr-wp.pdf

F5 Solution Center.
http://www.f5.com/solution-center/

F5.com.
http://www.f5.com

BIG-IP Global Traffic Manager.
http://www.f5.com/pdf/products/big-ip-global-traffic-manager-ds.pdf

BIG-IP Link Controller.
http://www.f5.com/pdf/products/big-ip-link-controller-ds.pdf

BIG-IP Application Security Module.
http://www.f5.com/pdf/products/big-ip-application-security-manager-ds.pdf

iRules.
http://devcentral.f5.com/Default.aspx?tabid=75

F5 Dev Central
http://devcentral.f5.com

Other References:

Sdforum press release
http://www.sdforum.org/SDForum/Assets/PDFs/Newsletters/Final_SDFnews_FebMar07.pdf