WHITE PAPER

Running Microsoft® Office SharePoint Server 2007 On VMware® Infrastructure

Virtualizing MOSS 2007

April 2009

vmware®

Contents

Executive Summary

Microsoft Office SharePoint Server (MOSS) 2007 is one of Microsoft's fastest growing server products. MOSS 2007 provides a highly-integrated applications platform that hosts a number of different collaboration and document management services. VMware and its partners have done significant testing to evaluate behavior and develop best practices for virtualizing MOSS deployments on VMware Infrastructure (VI). Our experience shows that all tiers in complex three-tier MOSS 2007 solutions can be successfully virtualized using VMware Infrastructure:

- MOSS 2007 scales well in virtual architectures, with more flexibility in design and deployment than in purely physical implementations.
- MOSS 2007 performance in a virtualized infrastructure is on par with performance in physical server deployments.
- Virtualizing the SQL Server database required for MOSS operation provides the advantages of higher availability and more flexible deployment without performance penalties, when basic best practices for system design are followed.
- The core benefits of VMware Infrastructure: higher consolidation, built-in availability, rapid provisioning, rapid application development lifecycles, enhanced security and dynamic scalability are all readily leveraged in a MOSS deployment on VMware Infrastructure.

In this document, we describe the implementation and best practices for virtualizing a typical medium-sized MOSS 2007 intranet farm deployment. We also have included results from tests run to measure and validate Web Front End (WFE), Query, Index, and SQL Server performance for this solution. While the solution described here is for a medium-size MOSS 2007 farm deployment, tests were also performed to validate both scale-up and scale-out techniques that could be used to support much larger MOSS farm deployments. This data would be useful in the design of a modular MOSS solution on a much larger scale.

Introduction

The purpose of this whitepaper is to provide guidance and best practices for implementing a mid-sized MOSS 2007 farm solution and to demonstrate how a VMware Infrastructure can be applied to enterprise SharePoint deployments of all sizes. In addition, we have provided references at the end of this document to other resources and publications detailing VI-based virtualized solutions for large-scale MOSS 2007 farm deployments.

MOSS 2007 provides enterprise-scale capabilities to manage content and business processes and simplify how people can find and share information. Using the combined features of Windows SharePoint Services and MOSS 2007, plus the design and customization capabilities of Office SharePoint Designer 2007, organizations can enable their users to create, manage, and easily build their own SharePoint sites, and make these sites more easily discoverable throughout the organization.

MOSS 2007 relies on the Windows SharePoint Services 3.0 technology to provide a consistent, familiar framework for lists and libraries, site administration, and site customization. All Windows SharePoint Services 3.0 features are also available in MOSS 2007. MOSS 2007 offers additional features that are not available on a Windows SharePoint Services site, including additional site templates related to enterprise and publishing scenarios.

MOSS 2007 farm implementations follow the three-tier web application architectural model, which requires multiple servers. Each server runs specific roles and supports various scale-up and scale-out options. The complexities of multiple server design, provisioning, resource management, and the often

mission-critical nature of the services provided, make MOSS 2007 systems excellent candidates for virtualization.

Despite these advantages, MOSS 2007 virtualization may not be a familiar subject to some experienced SharePoint system architects. VMware often gets questions about MOSS virtualization, including:

- *Has anybody else done it?*

  Yes, VMware customers have been virtualizing SharePoint deployments for years. Their objectives range from consolidation, to increasing availability, to ensuring that solution designs can easily grow and evolve.

  VMware Infrastructure has also been widely adopted by IT managers in data centers for the virtualization of other large scale applications including SAP, Exchange, SQL Server, and Oracle databases.

- *Will the SQL back-end run in virtual machines?*

  Yes, SQL runs great and can dramatically benefit from virtualization on the VMware Infrastructure platform. (Reference documents on SQL virtualization are noted at the end of this paper.) Following good design principles allow most SQL applications to run on virtual platforms without operational performance impact.

- *What if we choose not to virtualize SQL?*

  Deploying any multi-tier / multi-server architecture allows choices. You don't have to virtualize your entire SharePoint infrastructure all at once. Some customers start by virtualizing front-end and middle tiers first and then transition to an all-virtual infrastructure over time.

- *How do I control the risk of deploying on a new platform?*

  Deploying complex applications is easier in a VMware Infrastructure development environment. With virtual test bed and isolated virtual network, you can quickly and reliably test components in real-world scenarios. You can test components prior to production deployment and implement virtual components as you gain familiarity with the capabilities of the VMware Infrastructure platform. You will probably find that the benefits of virtualization far outweigh any risks.

- *Does Microsoft support MOSS 2007 deployed on VMware?*

  Absolutely. For customers with a Premier Support contract, Microsoft will support MOSS 2007 running in production. For other customers, Microsoft's Server Virtualization Validation Program (SVVP) provides support for qualified virtualization products. VMware products are all qualified for SVVP within 60 days of their release. For more information on Microsoft Support, go to the web site http://support.microsoft.com/kb/897615.

- *How can I meet my availability requirements if I virtualize MOSS 2007?*

  You can use the same techniques in a virtual environment, as are available in a physical infrastructure environment, for example, using network load balancing and clustering. In addition, the VMware Infrastructure environment provides additional high availability, dynamic resource allocation, load balancing and migration tools with VMware HA, DRS, and VMotion.

VMware Infrastructure has been widely adopted by all companies in the Fortune 500, and 85% of the global 1000, to deliver more reliable and flexible IT operations. VMware Infrastructure is a superior platform for running large scale applications and it has a proven success track record at running all sizes of deployments of business-critical applications including SAP, Exchange, SharePoint, SQL Server, and Oracle. You can read more about customer successes with VMware Infrastructure by visiting the site http://www.vmware.com/solutions/business-critical-apps/.

Business Challenges

A SharePoint infrastructure is frequently deployed first as a collaboration and document sharing tool. SharePoint makes workgroups more efficient and reduces the cost of coordinated collaboration, while allowing the flexibility needed by workgroups with diverse needs and goals. This flexibility and utility often leads to rapid growth in demand for both capacity and bandwidth as more users leverage these tools to coordinate workflows and manage more documents. More sophisticated SharePoint deployments can optimize business workflow and communications, and can quickly become critical components in everyday commercial operations.

While SharePoint offers very large benefits to organizations, these benefits are only realized when the underlying systems are available and perform at an acceptable level. Rapid growth and high availability are difficult features to manage in a traditional IT environment. Accommodating both often requires the high cost of over-designing and over-building at the earliest stages of deployment.

Since SharePoint encourages rapid growth and "viral" proliferation, user goals may conflict with the ability of the IT staff to deliver the services desired when needed within budgetary and manpower constraints. Flexibility is extremely valuable during this early period. If rapid growth and evolution can be supported at realistic costs, SharePoint can become an important tool to rapidly increase everyday productivity.

VMware Infrastructure can facilitate this capability, allowing organizations to leverage the benefits of SharePoint on a "pay as you go" basis. Since high availability features are inherent to VMware Infrastructure products, these can be leveraged as needed, on demand. By virtualizing SharePoint, the classic challenges of deploying a complex, high-growth IT service are alleviated, allowing resources to be spent on maximizing the value of the tool in everyday business practice.

Contrast the benefits of a virtual infrastructure with the limitations of a traditional deployment. Using conventional physical infrastructure typically leads to over-provisioning. This creates significant resource underutilization and high system power, cooling, and operating costs. In addition, for complex architectures such as MOSS 2007, using physical servers and infrastructure may have other limitations:

- Application delivery is traditionally gated by the need for manual configuration and provisioning for each new application or configuration change on a specific hardware platform. This can be slow and, in an existing infrastructure, lead to excessive downtime. It can also constrain growth to the organization's ability to purchase new hardware.

  Virtual deployments typically take minutes, can share currently deployed hardware, and can be adjusted "on the fly" when more resources are required.

- Application architectures, such as those provided by MOSS 2007, are rapidly evolving towards highly distributed, loosely-coupled applications. The conventional x86 computing model, in which applications are tightly coupled to physical servers, is too static and restrictive to efficiently support these complex applications.

  With a virtual deployment, the architecture can be as modular as is appropriate, without expanding the hardware footprint. The dynamic nature of virtual machines means that the design can grow and adapt as required, and the need for a "perfect" initial design is eliminated.

- Availability becomes a critical factor. In a physical environment, the cost and complexity of server clustering is required to increase availability. In a highly distributed environment, VMware Infrastructure can increase application availability at a much lower cost than using traditional HA strategies.

- Rising data center costs (for power, cooling, floor space, etc.) even while some server computing resources go under-utilized. It is well understood that server consolidation through virtualization is a

significant factor in reducing cost. The higher density available from VMware Infrastructure versus other products can provide significant savings on hardware and Microsoft licenses.

This paper is intended to demonstrate that virtualization with VMware Infrastructure can minimize the challenges in a SharePoint deployment, to ensure that maximum commercial value is realized from an investment in SharePoint systems.

Benefits of MOSS 2007 Virtualization

VMware Infrastructure is not only well suited to handle the demands of MOSS 2007 deployments; it also provides a superior platform and brings significant benefits to any multi-tier infrastructure. With VMware Infrastructure, you can deliver your business-critical applications as dynamic, cost-efficient, and reliable IT services, scaling applications dynamically and providing built-in availability and simple disaster recovery. Here are some of the benefits that applications such as MOSS 2007 can take advantage of by using VMware Infrastructure:

- **Higher Availability for components and services.** By leveraging the inherent benefits of a virtualization-based platform, VMware Infrastructure offers a variety of availability options. Each of these options provides different levels of both protection and cost, capable of meeting the unique high availability requirements of any organization. A number of tools available from VMware, Microsoft, and third-party software and hardware vendors can be used to facilitate higher operating availability under normal circumstances, as well as provide on-site and remote site availability and disaster recovery (DR).
- **Guarantee application Quality of Service.** Ensure end-user QoS by automatically providing the right levels of application availability and scalability. Application availability and scalability levels can be dialed up or down dynamically as business requirements evolve, so you can meet QoS requirements in the most cost-effective way.
- **Application delivery as dynamic, cost-efficient, and reliable IT services.** VMware offers the most proven, high performance, and reliable virtualization platform that is ideally suited to running modern distributed applications. With VMware Infrastructure, an IT staff can deliver business-critical applications as dynamic, cost-efficient, and reliable IT services on the internal cloud, free from the constraints of static, dedicated infrastructure.
- **Performance and reliability that can often exceed the performance of physical servers.** VMware Infrastructure, running applications including large databases, enterprise business suites, and email, can be easily scaled on the fly—to take advantage of multi-core servers, higher network bandwidth, and enterprise storage, and deliver outstanding IO, CPU and memory performance. Applications can also be scaled out in multiple virtual machines to increase the throughput achieved per physical host.
- **Minimize infrastructure costs through higher consolidation ratios for demanding applications.** Denser consolidation enables significantly lower overall costs for hardware, Microsoft software licenses, and provisioning for networking and storage. Higher densities also eliminate the need to provision dedicated testing, staging, training, and DR servers, since isolated virtual networks can be created to test software in secure virtual machines. In production, VMware Infrastructure enables applications to use exactly the capacity they need, when they need it, and eliminate the need to overprovision capacity.
- **Accelerate application delivery.** Create pre-configured templates and virtual appliances that can be provisioned on-demand; test multi-tier applications quickly and efficiently; and automate release cycles. Deploy standard, pre-configured application stacks, ensuring consistency across production applications and minimizing manual configuration overhead and configuration errors. Also, streamline testing and troubleshooting with VMware snapshots and clones.

- **Leverage broad ISV ecosystem.** Three of the top four global ISVs – Microsoft, SAP, and IBM – now provide the same level of application support on VMware Infrastructure as on physical servers. Hundreds of smaller ISVs support their applications on VMware Infrastructure, recognizing the value that VMware brings to their customers.

MOSS 2007 Virtualization Solution

The main goals of the solution described in this document were to research the behavior of a virtualized mid-size MOSS 2007 farm solution, develop and provide best practices to customers, and validate that performance of MOSS 2007 running in a VMware Infrastructure environment was on-par with native (non-virtualized) deployments.

The MOSS 2007 farm solution described in this document virtualized an original physical system hardware design recommended by the Microsoft System Center Capacity Planner using system workload and other requirements described later in this section. As such, the solution described in this document did not take advantage of additional benefits that might have been derived by including VMware DRS, HA, and VMotion. Documents referenced at the end of this paper describe solution architectures and designs for other MOSS 2007 deployment scenarios, including very large system implementations and highly scaled solutions including benefits from VMware DRS, HA, and VMotion.

Characterizing a Mid-Size MOSS 2007 Farm Solution

MOSS 2007 is an integrated suite of server capabilities that can help improve organizational effectiveness by providing comprehensive content management and enterprise search, accelerating shared business processes, and facilitating information-sharing across boundaries such as business and Intranet networks, and the Internet. MOSS 2007 supports all intranet, extranet, and Web applications across an enterprise within one integrated platform, instead of relying on separate fragmented systems.

Organizations can use MOSS 2007 to facilitate collaboration, provide content management features, implement business processes, and supply access to information that is essential to organizational goals and processes. Here are some of the features and capabilities that are commonly included in MOSS 2007 deployments:

- **Collaboration** – Allow teams to work together effectively, collaborate on and publish documents, maintain task lists, implement workflows, and share information through the use of wikis and blogs.
- **Portals** – Create a personal MySite portal to share information with others and personalize the user experience and content of an enterprise Web site based on the user's profile.
- **Enterprise Search** – Quickly and easily find people, expertise, and content in business applications.
- **Enterprise Content Management** – Create and manage documents, records, and Web content.
- **Business Process and Forms** – Create workflows and electronic forms to automate and streamline your business processes.
- **Business Intelligence** – Allow information workers to easily access critical business information, analyze and view data, and publish reports to make more informed decisions.

In addition, MOSS 2007 deployments can incorporate any features that are available in Windows SharePoint Services 3.0, which is an enabling technology that is included in Microsoft Windows Server. While organizations can completely customize the features provided by a specific MOSS 2007 deployment, Microsoft has also included a number of different "out-of-box" site templates that bundle SharePoint service features and capabilities for common application by enterprises. For example, two of the site template categories are:

- **Publishing** – design, review/approval, and posting of web content.
- **Collaboration** – document-centric features that incorporates business processes for creating, publishing, and sharing Intranet web content and documents.

The solution described in this paper is based on a medium-size MOSS 2007 farm deployment that uses the Intranet Collaboration template-based mix of different services. To get a baseline recommendation, we used the Microsoft System Center Capacity Planner (http://technet.microsoft.com/en-us/library/bb961988.aspx) and specified criteria for the solution to meet the following requirements: 1000 Users, Average Collaboration, Intranet Portal, High Availability, and 250GB available for content storage.

Figure 1 shows a topology diagram generated by the Microsoft Capacity Planner for the mid-sized MOSS 2007 farm solution. The MOSS Farm solution represented in this illustration shows a deployment design created for hosting a Corporate Intranet Portal site using the SharePoint Collaboration Template. The farm includes a total of five servers: two load-balanced Front End Web Servers (which also hosts Query Servers), an Index Server, and Microsoft SQL Server Cluster (active/passive) to host a SQL Server database used by the solution.

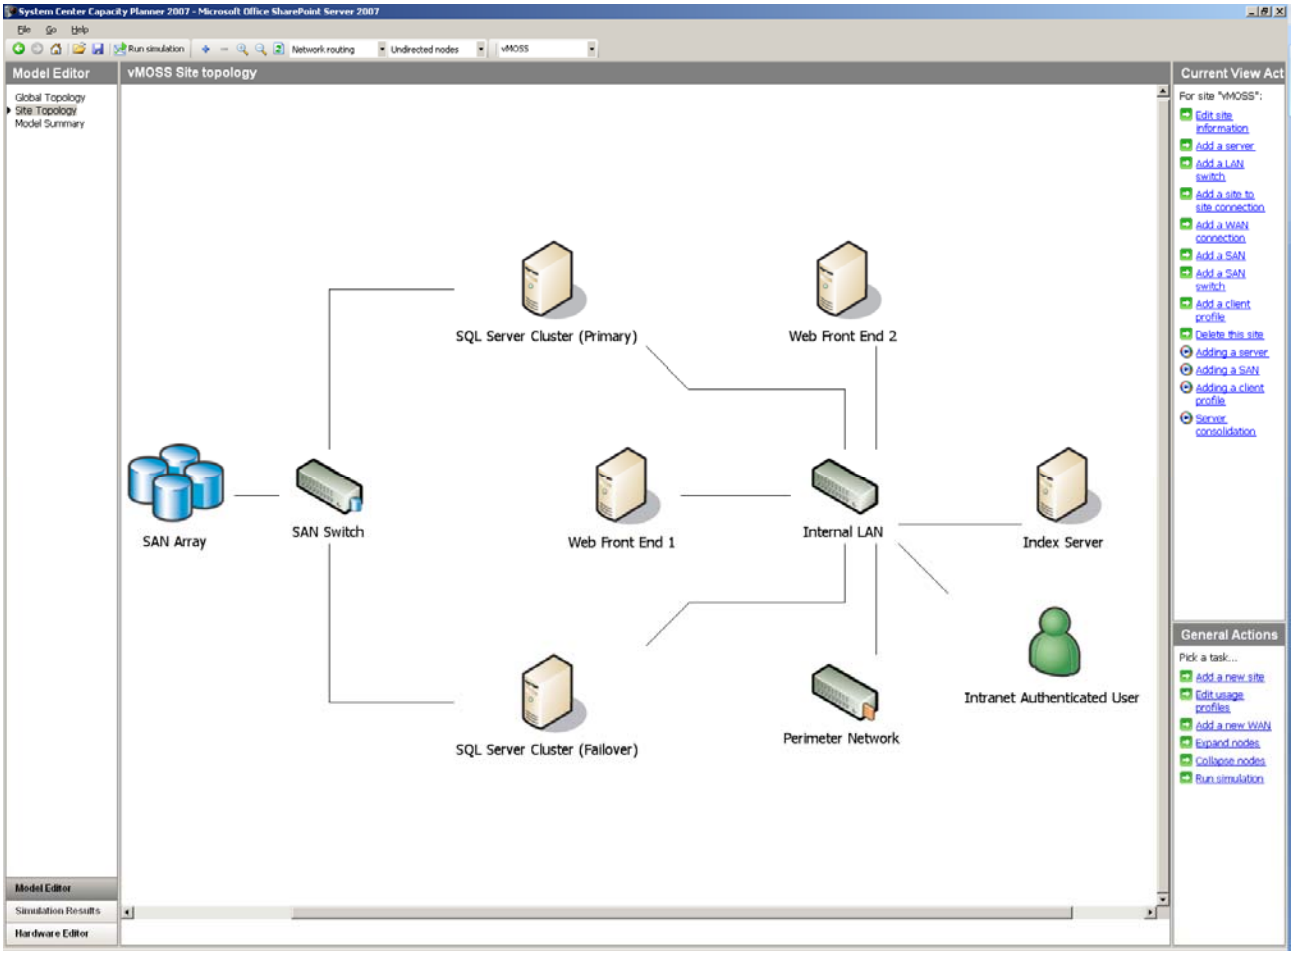

Figure 1. MOSS 2007 Farm Site Topology Diagram Created by Microsoft Capacity Planner

Figure 2 shows the logical architecture design translation for the MOSS 2007 mid-sized farm deployment.

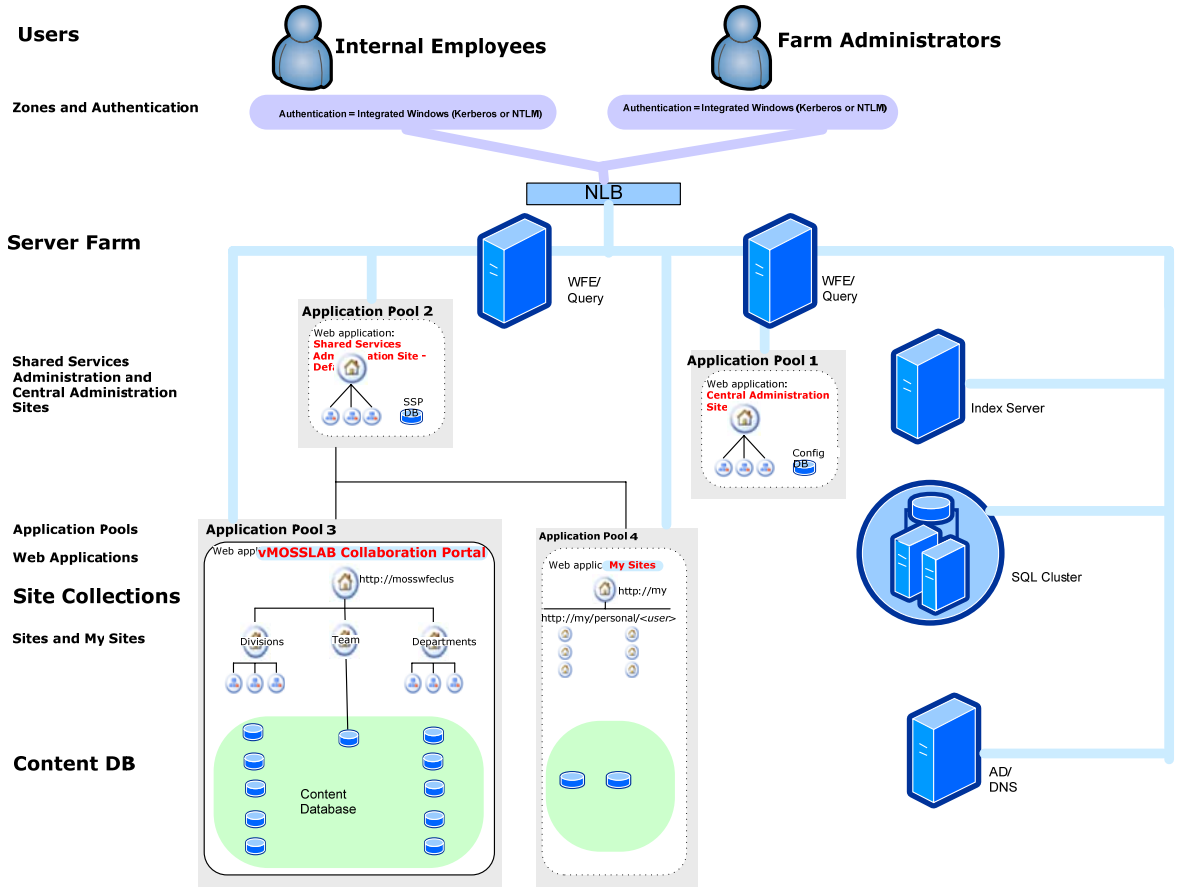

Figure 2. MOSS 2007 Farm Logical Architecture

The figure illustrates the following key aspects of the MOSS 2007 mid-sized farm deployment:

- Integrated Windows authentication is used to address the security requirements of the Portal.
- A Shared Service Provider is used to provide common services such as Search and User Profiles.
- Multiple application pools are used to achieve process isolation between content.
- Multiple Web Applications are used to isolate users, enforce permissions, and optimize performance and manageability.
- Multiple Site Collections are used to bridge the logical architecture with information architecture.
- Multiple Content Databases are used to provide better content and database manageability.

Virtualizing the Mid-Size Moss 2007 Farm Solution

A basic premise in converting the physical, non-virtualized mid-size MOSS 2007 farm solution to a virtualized design was that AD/DNS (and other necessary infrastructure) already existed and the MOSS 2007 Farm solution is only added incrementally to the existing environment. In the virtualized solution, the MOSS farm is configured across two physical servers to support hardware redundancy. Server 1 is configured with Web Front End and the active node of SQL Cluster. Server 2 is configured with Web Front End, Index Server, and the passive node of SQL Cluster. Windows Network Load Balancing is used for Web Front End failover and MSCS (Microsoft Clustering Service) is used for database failover.

Figure 3 shows the virtualized implementation of the MOSS 2007 farm solution.

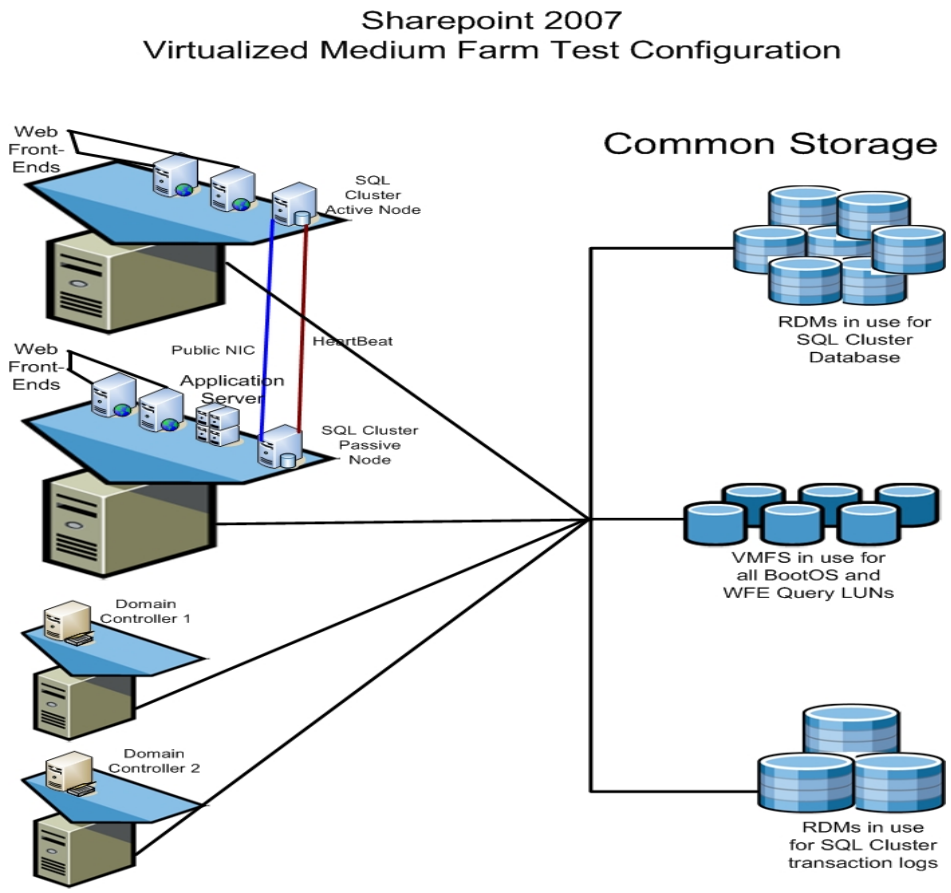

Figure 3. Virtualized MOSS 2007 Farm Solution

The virtualized solution was created using VMware ESX 3.5 u2 running on two HP BL460c Blade servers. The solution used 64-bit SharePoint Server 2007 SP1, SQL Server 2005 SP3 running on Windows Server 2003 SP2, Microsoft Network Load Balancing and Active/Passive Microsoft Cluster Server (MSCS) software. WFE Instances and the Index Server also were run on Windows Server 2003 SP2. The SharePoint Server 2007 Root Portal site was set up to use the Collaboration template.

The workload set up in the virtualized MOSS 2007 farm solution included the following:

- 260 GB content – consisting of approximately 600,000 Items; the index server was set up for weekly full crawl, with incremental crawl performed every 4 hours.
- 10 Site Collections – 100,000 documents with unique metadata spread equally across ten site collections and 250 Team Sites with approximately 500,000 documents and list items spread across seven site collections.

The physical infrastructure for the virtualized solution consisted of two GB Ethernet switches, two redundant paths to storage using a Fiber Channel 8Gb Fabric switch utilizing 4Gb ports, teamed NICs, and the two HP BL460c blade servers, which are dual socket quad core processor systems.

Storage was provided by an EMC Celera NS42 array with two types of storage:

- iSCSI volumes were used for the boot OS LUNs and Query Catalog.
- RDMs were used for the SQL Database and SQL Cluster transaction logs. (RDMs were required since we were following best practices for Microsoft Clustering.) Please review the *VMware MSCS Setup Guide*, available at the following VMware web site location, when implementing Microsoft Clustering on VMware Infrastructure 3:

  http://www.vmware.com/pdf/vi3_35/esx_3/vi3_35_25_u1_mscs.pdf

Figure 4 shows the network and physical infrastructure of the virtualized MOSS 2007 farm solution.

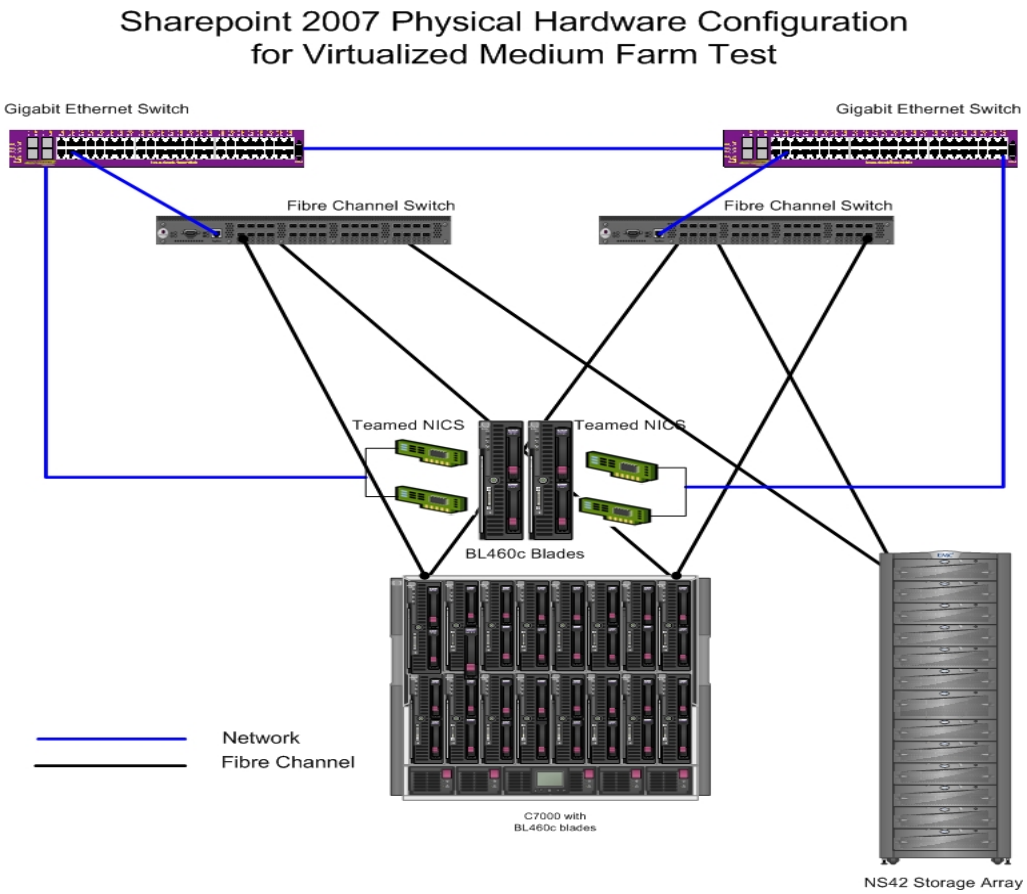

Figure 4. Network and Physical Infrastructure Configuration of the Virtualized MOSS 2007 Farm Solution

Solution Configuration Steps

This section highlights the key steps involved in virtualizing the MOSS 2007 farm solution on VMware Infrastructure:

1. Use the Microsoft SharePoint Capacity Planner to get a sample data point.

   (http://technet.microsoft.com/en-us/library/bb961988.aspx)

   This tool takes numerous inputs of work load parameters and requirements you specify and then outputs a recommended physical configuration that you can use as a data point.

2. Plan for administrative and service accounts as described in the Microsoft TechNet SharePoint Server TechCenter article "Plan for administrative and service accounts" (http://technet.microsoft.com/en-us/library/cc263445.aspx).

3. Create VM templates that you can reuse to create additional instances of Web Server, Index Server, or Application Server required in the farm. Settings in the VM templates would typically include the following:

   - 64-bit Windows Server (2003 SP2 or 2008). The 64-bit version is recommended since all future versions of MOSS will be 64-bit.
   - Use two or more vCPUs in the template to support SMP and HAL.
   - Configure IIS, install .NET Framework 3.0, and enable ASP.NET 2.0 as detailed in the following Microsoft TechNet SharePoint Server TechCenter article that provides information for deploying a simple server farm:

     (http://technet.microsoft.com/en-us/library/cc262243.aspx).

   - Use a separate vSCSI adapter for the Query LUN to store the Query Catalog (size of the catalog will vary in proportion to amount of indexable content; 5 to 10% is typical).
   - Allocate 16-32 GB drive for Boot operating system.
   - Use VMFS for maximum flexibility.

4. For Web Front End Server configuration, create an instance from the VM template described in Step 3, with the following settings:

   - Start with 2 vCPU/4GB RAM and scale up or out depending on your requirements.
   - Use separate controllers.
   - Add disk before configuring the network settings.
   - Use the enhanced vmxnet adapter.
   - Ensure you have enough space for SharePoint and IIS Logging.
   - Configure your preferred hardware or software based network load balancer (NLB). If you are configuring Windows Network Load Balancer, refer to instructions provided in the VMware article located at:

     http://www.vmware.com/files/pdf/implmenting_ms_network_load_balancing.pdf

5. For the Index Server configuration, create an instance from the VM template described in Step 3, with the following settings:
   - Start with 2 vCPU / 6GB RAM and scale up to 4 vCPU / 8GB depending on your requirements.
   - Use separate controllers.
   - Add disk before configuring the network settings.
   - Use the enhanced vmxnet adapter.
   - Ensure you have enough space for SharePoint and IIS Logging.
   - If possible, use dedicated WFE for crawling contents.
   - To minimize network traffic, consider placing virtual machines on the same VMware ESX host to perform networking operations within VMkernel.
6. For SQL Server configuration, create a new virtual instance with the following:
   - 64-bit Windows Server (2003 SP2).
   - Start with 2 vCPU / 6GB RAM and scale up to 4 vCPU / 12GB depending on your requirements.
   - Use RDMs. (When using Microsoft clustering to cluster SQL database, we used RDMs to follow recommended best practices; if the database tier is not clustered using MSCS, VMFS is a better option.)
   - Use eagerthickzeroed disks. (The eagerthickzeroed format provides a full format that wipes the entire disk to make sure no old data is present; also plays a role when setting up clustering.)
   - A separate vSCSI controller must be used for OS and clustered data disks.
   - Use separate vSwitches for Public and Heartbeat.
   - Team Physical NICs.
   - Install and configure SQL 2005 on MSCS cluster as detailed at:

     http://www.vmware.com/pdf/vi3_35/esx_3/vi3_35_25_u1_mscs.pdf
7. Install MOSS 2007 SP1 on the WFE/Index/Application server instance. If possible, automate SharePoint installation and configuration using scripts for setup.exe, psconfig.exe and stsadm.exe options.

As mentioned earlier, the MOSS 2007 portal site described in this solution is configured with the Collaboration template for the root portal. In addition, the solution also included the following configuration steps or settings:

- The ten site collections described in the solution were hosted in a separate database using "stsadm -o createsiteinnewdb" options as detailed at the following location:

  http://technet.microsoft.com/en-us/library/cc262407.aspx

- Scripts were used to relocate the new SharePoint Content DB and LOGS from their default locations.
- The SharePoint 2007 Test Data Population Tool was used to populate approximately 600K items with a total size of approximately 260 GB across the site collections. This tool is available at:

  http://sptdatapop.codeplex.com/Release/ProjectReleases.aspx?ReleaseId=1141.

- The Index Server was scheduled for a weekly full crawl and an incremental crawl every 4 hours.
- IIS Logging was enabled since most enterprises use analytical tools like WebTrends (which depends on IIS Log).
- SharePoint Server 2007 was configured with default settings. To represent real-world scenarios, no additional caching or web garden features were enabled.

Some other general tips and recommendations for virtual MOSS 2007 solution deployment include the following:

- Install any critical and recommended Microsoft hotfixes.
- To enhance manageability and performance of SharePoint, limit the content database to 100 GB.
- Install the Central Administration Web application on the Application Server if you have one; otherwise, configure on the Web Front End Servers.
- Configure the Search Crawl account with Read Only access so that only published documents are indexed. Ensure the content crawl schedule does not overlap with peak usage and down times scheduled for your environment.
- Ensure adequate storage space is allocated for the SharePoint Content Index file based on guidelines provided in the Microsoft TechNet article located at http://technet.microsoft.com/en-us/library/cc262574.aspx.
- Post configuration – use the Microsoft Best Practice Analyser for SharePoint located at http://www.microsoft.com/downloads/details.aspx?familyid=cb944b27-9d6b-4a1f-b3e1-778efda07df8&displaylang=en to ensure all recommended best practice rules have been applied to your solution configuration.

Scaling and Performance Test Results

VMware extensively tested the virtual MOSS 2007 farm solution described in this document. In particular, tests were run to compare how WFE servers in the MOSS 2007 farm solution performed against un-virtualized physical WFE servers. In addition, we performed tests to see how the virtualized solution might benefit from various scaling methods and the performance improvements or increases customers would likely realize.

Test Methodology

Since there are not any clear standard benchmark tests and procedures as yet established for measuring the performance of SharePoint deployments, we primarily followed Microsoft's recommendations for methods to capture key statistics and obtain measures and indicators of solution performance and test results. These included:

- Following the guidelines for estimating performance and capacity requirements for Office SharePoint Server as detailed in the Microsoft TechNet SharePoint Server TechCenter article located at:

  http://technet.microsoft.com/en-us/library/cc261795.aspx

- Obtaining the User Profile portal statistics from IIS Logs/ WebTrends/ SharePoint Usage Statistics/ Others. If tools such as those provided by WebTrends are not available at your site, you can use CodePlex SharePoint test data load tools located at:

  http://www.codeplex.com/sptdatapop/Release/ProjectReleases.aspx?ReleaseId=1214

  and

  http://www.codeplex.com/sptdatapop/Release/ProjectReleases.aspx?ReleaseId=1141

- Using Visual Studio Team System Test Edition, run CodePlex WSSTestProject and MOSSTestProject tests to capture Request per Second (RPS), Avg. Response Time, and other key performance counters.
- In addition, using VMware ESX, you can use the following command to capture ESX counters on ESX hosts:

  ```
  esxtop –a –b
  ```

Specific tests of the MOSS 2007 farm solution described in this guide (for both physical and virtual implementations) were performed to measure the following characteristics and impacts on performance:

- Adding multiple Web Front Ends
- Adding additional vCPUs
- User Experience while Full Text Crawl is on.

The following workload was used in creating benchmarks for specific tests:

- All tests used an 80-10-10 profile Mix of Read-Write-Search.
- Tests measured four different Read actions in CodePlex WSSReadMix (WSSHomePage, WSSDispForm, WSSSmallDocLibAllItems, WSSDocHTTPFetch) applying equal weight to each.
- Tests use one Write action (WSSPutDocumentCoded) from CodePlex WSSReadWriteMix.
- A custom web test was created for Search and plugged in to our Load Test Mix.
  - Both MetaData Search and Free Text Search behavior was tested by changing the Search Mix in each load test. Metadata searches are performed almost entirely from the index catalog of the Query Server which is hosted with WFE servers in the MOSS farm, whereas free text searches typically involve data stored in the WFE/Query Server cache, as well as data retrieved from SQL databases.
  - The MetaData Search web test was performed using a unique Title name randomly picked from a set of 100K documents.
  - The FreeText Search web test was performed using a keyword that returned between 5 and 10 documents, on average.
- A Visual Studio Team System (VSTS) Load Generator, with zero think time, was used throughout the tests.

Test Results

In this section, we describe and list the results from individual tests, explain specific data results, and describe likely conclusions that can be drawn from each result.

Physical versus Virtual WFE Server Comparison

Good WFE performance is key to the overall performance of MOSS 2007 farms. The following chart shows results comparing relative WFE request per second (RPS) of physical WFE servers versus a virtual WFE server utilizing either one or two processors.

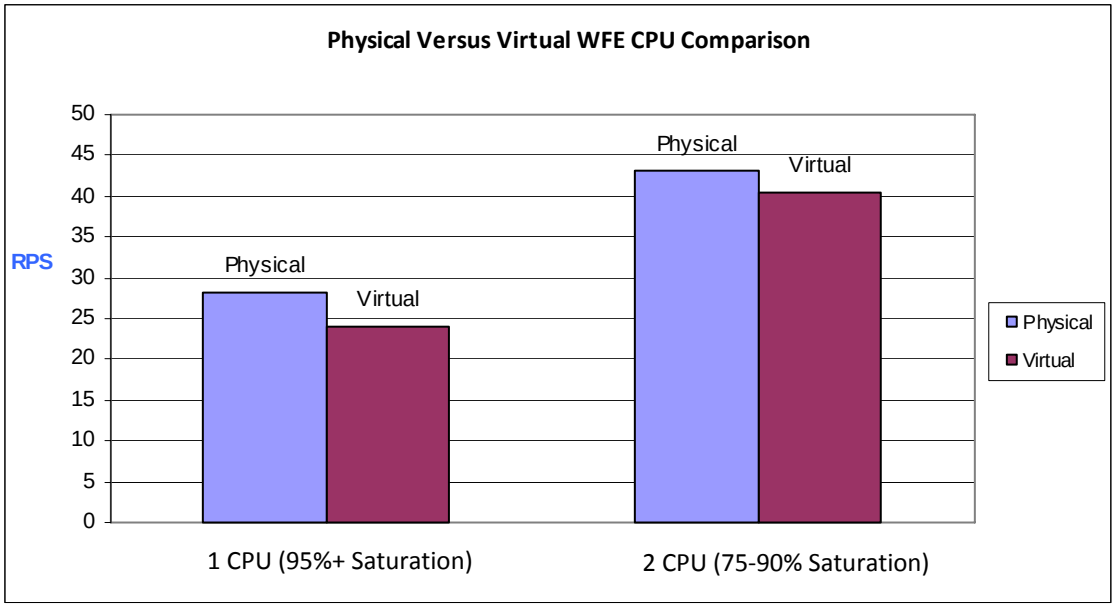

From the test results, we can see that overall RPS differs very little between physical and virtual WFE servers, even at higher CPU saturation levels.

One vCPU versus Two vCPU Server Comparison

The following test was used to determine whether increasing the number of CPUs per server in a virtual environment had any significant impact on WFE performance. The first set of columns establishes a data point that compares WFE requests per second (RPS) for four WFE servers (with one CPU per server) versus two WFE servers (with two CPUs per server) for the same workload. The results from this data point show comparable performance, with 1-vCPU servers performing slightly better than 2-vCPU servers. The second set of columns provides a second data point comparing performance of six WFE servers (with one CPU per server) versus three WFE servers (with two CPUs per server for the same workload). Again, the results showed slightly better performance for 1-vCPU servers versus 2-vCPU servers.

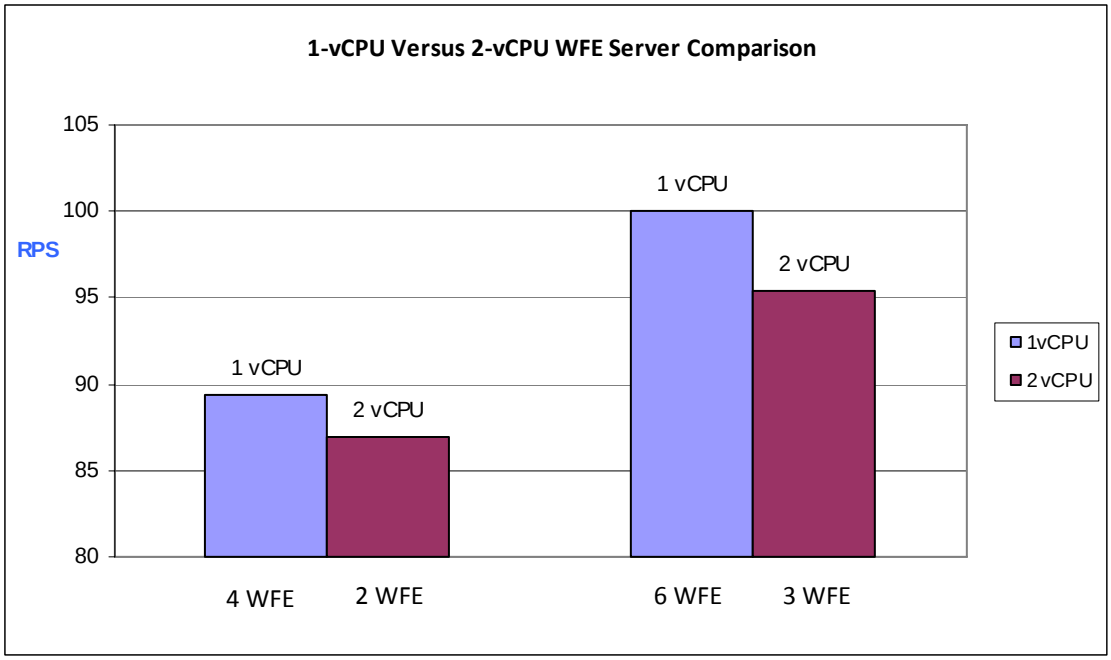

Note that although the 1-VCPU servers provides slightly better performance, other considerations such as MOSS 2007 licensing costs per server may determine the most suitable configuration for specific customer environments.

80-10-10 Mix, Constant Load, Free Text Search using 1vCPU WFE

The following test with an 80-10-10 user profile mix of Read - Write- Free Text Search was performed to determine how well SharePoint deployments would scale with additional virtual WFE servers being added to support searches. We also wanted to see if performance would increase linearly with the expansion to use more servers.

The test was carried out with a constant load performing concurrent random free text searches. Both RPS and response times were measured for different configurations ranging from one to five virtual WFE servers.

Here is a chart of the performance results tracking RPS and average response time measurements over the range of one to five virtual WFE servers.

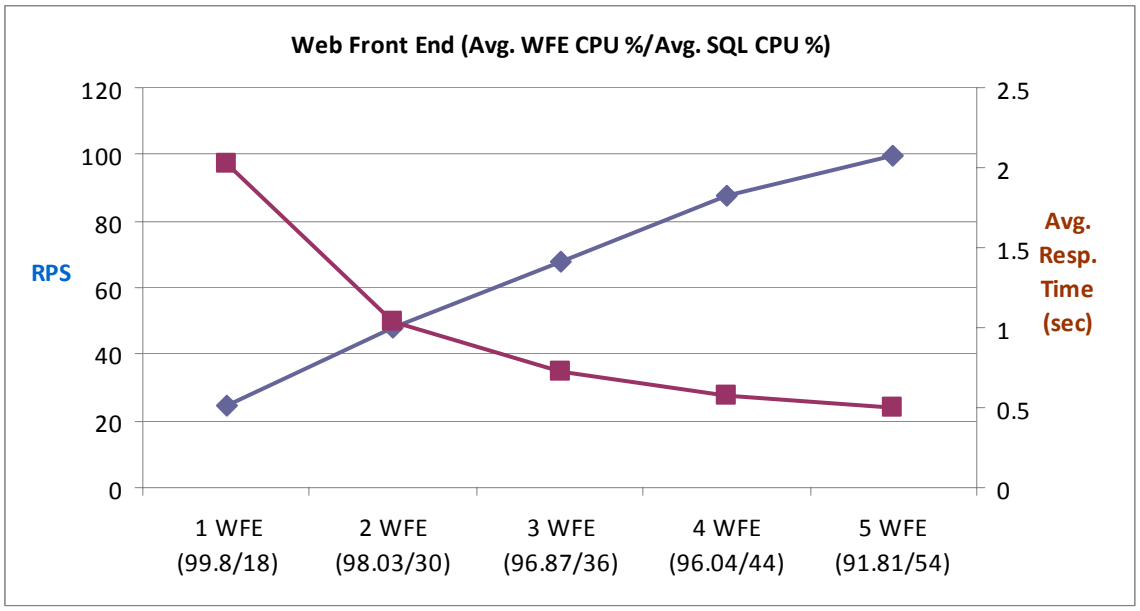

As the rising blue (RPS) line indicates, throughput increase was mostly linear with RPS increasing proportionally with the addition of each new WFE server. In addition, the average response time also improved with each new server added. The SQL Server CPU utilization displayed in the X axis label (below each WFE callout) shows the virtualized SQL Server cluster supporting additional web front ends and increased throughput.

80-10-10 Mix, Constant Load, MetaData Search using 1vCPU WFE

The following test was similar to the previous one, except that scaling performance was measured with metadata searches instead of free text searches. (Metadata searches are performed almost entirely from the index catalog of the Query Server which is hosted with WFE servers in the MOSS farm, whereas free text searches typically involve data stored in the WFE/Query Server cache, as well as data retrieved from SQL databases.)

As in the free text search test, the metadata search test was carried out with a constant load. Both RPS and response were measured for different configurations ranging from one to five virtual WFE servers.

Here is a chart of the performance results tracking RPS and average response time measurements over the range of one to five virtual WFE servers.

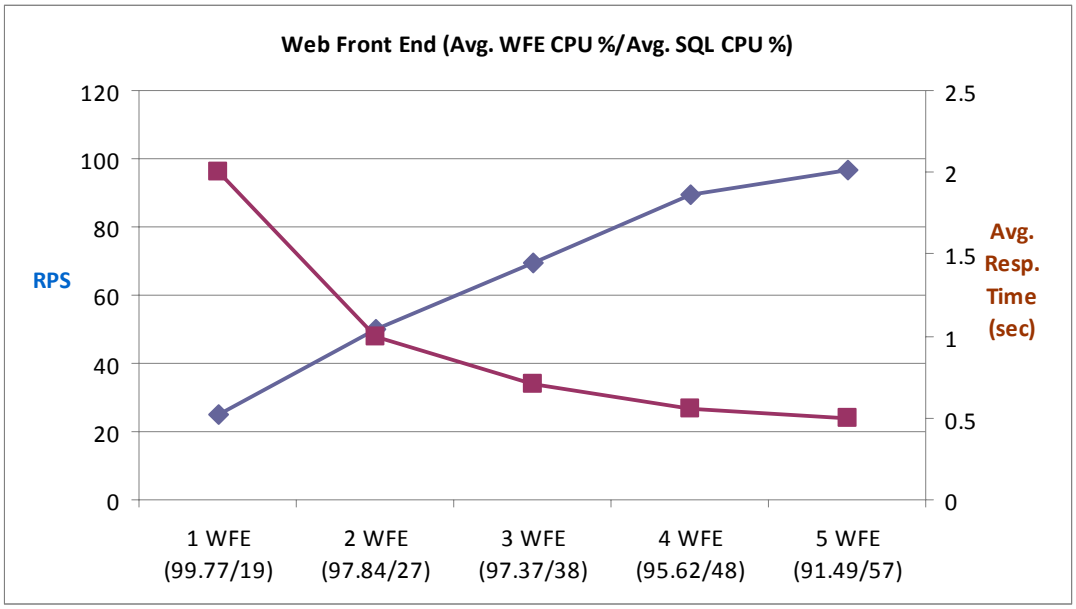

Similar to the free text search test, the RPS response to additional WFE servers again was mostly linear with RPS increasing proportionally with the addition of each new server. In addition, the average response time also improved with each new server added. The SQL Server CPU utilization displayed in the X axis label (below each WFE callout) shows the virtualized SQL Server cluster supporting additional web front ends and increased throughput.

80-10-10 Mix, MetaData Search, 75-85% WFE CPU using 1vCPU WFE

CPU utilization is one of the key metrics used to monitor the health of a MOSS farm. In the following test, we measured the throughput of the farm by varying the user load and keeping the WFE CPU in the 75-85% range to observe the performance at less than fully saturated (>90%) levels.

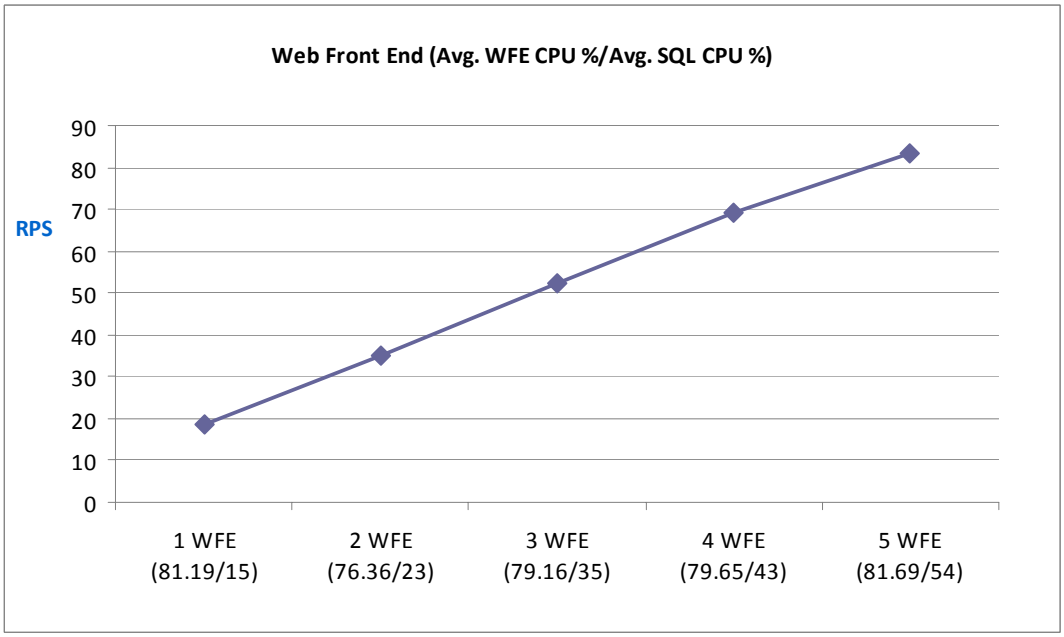

This test shows that as user loads were increased, bringing additional servers online helped maintain throughput and linear scaling.

Index Server Test

MOSS 2007 deployments require a schedule for Index crawling. Depending on the size of the overall content being stored, Index crawling windows (normally scheduled for off-peak hours) may overlap peak/production windows. A test was performed to measure the effect that index crawling may have on virtual WFE server response times and CPU utilization/saturation levels, to make sure that virtualized SharePoint systems still provide reasonable throughput and response times even during a full index crawl operation. The following table shows results for a system that includes six 1-vCPU WFE servers and compares average response times for metadata random access searches and CPU utilization levels during normal operation as well as times when a full index crawl is being performed.

Description	Average Response Time (in seconds)	Average WFE CPU Utilization	Average SQL CPU Utilization
80-10-10 Mix Test without index crawl	0.29	66.39 %	54 %
80-10-10 Mix Test with full index crawl	0.34	69.59 %	67 %

NOTE: The Full Index Crawl on 600K items (260GB) was completed in 8 hours using an average of 76% Index Server CPU utilization and 36% SQL Server CPU utilization.

These results indicate that overall Index and SQL utilization levels do go up during the index crawl and that the overall system throughput and response times undergo marginal degradation (20% or so). The effects reported are based on our test environment setup; it is important that you test and evaluate the effects of Index crawl in your specific environment.

Windows Server 2008 vs Windows 2003 Study

Windows Server 2003 is currently the most widely used operating system in production environments and this whitepaper is focused mainly on testing with Windows Server 2003. However, for the benefit of organizations that are evaluating or planning to move to Windows Server 2008 operating system, we did a similar study of MOSS farm virtualization on the Windows Server 2008 operating system. New Windows Server 2008 capabilities improve the management, security, reliability, and performance of Web applications like SharePoint. The new TCP/IP implementation in Windows Server 2008 contains optimization for high-latency networks and provides significant performance improvements over earlier versions of Windows networking.

For the Windows Server 2008 study, we built an identical MOSS Farm using Windows Server 2008 operating system across all tiers. The back-end database remained using SQL Server 2005, but used VMware HA for failover, instead of MSCS. Our study results, for the same Test Mix as performed with Windows Server 2003, show Windows Server 2008 having similar throughput and response times as Windows Server 2003 for 1vCPU WFE tests, and slightly better throughput and response times (in the range of 10-15%) for two or more vCPU WFE configurations.

Key Findings and Conclusions

Here are the key findings and conclusions that can be drawn from the tests measuring both physical and virtual SharePoint system performance:

- All tiers of MOSS 2007 system solutions can be virtualized and test results showed that the virtualized SharePoint deployments provided performance comparable to physical server deployments.
- Physical versus virtual WFE server tests showed very low virtualization overhead and minimal RPS throughput difference.
- 1-vCPU versus 2-vCPU WFE server tests showed 1-vCPU servers providing slightly better throughput results, however, other considerations such as licensing costs per MOSS 2007 server may determine the most suitable configuration for specific customer environments.
- WFE server scale-out tests expanding the number of virtualized servers from one to five (at non-saturated CPU levels) showed near linear performance; load tests for scalability tests with both Full Text and MetaData Search profiles provided comparable results.
- Load testing, with a Full Index Crawl running in the background, showed virtualized SharePoint system continued to provide reasonable response to end-user searches (approximately 15 to 20% drop in throughput/response time).
- SharePoint can cause significant spikes in SQL CPU utilization when executing more complex and expensive SQL queries, for example, with document modifications operations, check in/out, recursive list item queries, and Free Text Search with many hits – depending on the mix of such transactions in your overall (or planned) workload, SQL resource design will be an important consideration (including scale-out options for SQL Server databases within a MOSS 2007 deployment). Ensure that you set up a proper test environment that reflects your deployment requirements and evaluate your SQL design considerations accordingly.

Summary

In this paper, we described the process by which customers can approach the tasks of virtualizing current physical or un-virtualized MOSS 2007 solutions on VMware Infrastructure. We also compared scaling and performance tests for both physical and virtual solution deployments and those results showed that the performance of virtual solutions was on par with their physical server counterparts. However, customers who deploy virtual MOSS 2007 solutions with VMware Infrastructure stand to gain from other benefits that virtualization provides, such as server consolidation and higher utilization rates, which lowers overall operating and licensing costs, as well as application delivery as Dynamic, Cost-Efficient, and Reliable IT Services on your internal Cloud. Tests also indicate that VMware Infrastructure can make good use of both scale-up and scale-out techniques for MOSS 2007 deployments, which allow customers to build virtual infrastructure to support the needs of practically any size business or enterprise.

Here are some other specific conclusions we derived from the virtualization of the MOSS 2007 solution deployment described in this paper:

- All tiers in complex three-tier MOSS 2007 deployments can be successfully virtualized providing considerable benefits in keeping costs low and increasing solution flexibility and efficiency. Virtual overhead was minimal versus physical configurations. Performance and reliability matches and can even exceed the performance of physical servers running larger and more complex applications including large databases.
- Virtualized MOSS 2007 deployments scale well both scaling up and out for WFE and SQL tiers, maintaining high throughput and good response times. WFE scalability in virtual solutions was proven, maintaining performance and response times, even at high saturation levels. 1vCPU WFE results indicate a slight edge in performance, but using 2vCPU to scale WFE may provide a better approach to balance licensing costs and overall throughput.
- Virtualization allows for more rapid provisioning; VMware LAB Manager is also a good fit for SharePoint Application Lifecycle Management.

Although not directly demonstrated in this paper, other studies referenced in this paper also showed the benefits of applying VMotion, HA, DRS for larger scale MOSS 2007 deployments, resulting in more flexible resource allocation to provide higher availability, provide protection against planned downtime and host failure, and increase overall efficiency and throughput.

Additional Resources and Document References

Documents listed here provide additional information relevant to MOSS 2007 that is available on the VMware web site or from other third parties.

VMware Publications

- Virtualizing Microsoft SQL Server:

  http://www.vmware.com/solutions/business-critical-apps/sql/

- Whitepapers on VMware Performance Topic:

  http://www.vmware.com/resources/techresources/

- Setup for Microsoft Cluster Service on ESX:

  http://www.vmware.com/pdf/vi3_35/esx_3/vi3_35_25_u1_mscs.pdf

- Implementation of Microsoft Network Load Balancing in a Virtualized Environment:

  http://www.vmware.com/files/pdf/implmenting_ms_network_load_balancing.pdf

Microsoft Publications

- Microsoft Support Policy for software running in non-Microsoft hardware virtualization software:

  http://support.microsoft.com/default.aspx/kb/897615

- White papers for MOSS 2007:

  http://technet.microsoft.com/en-us/library/cc262733.aspx

- Estimate performance and capacity requirements for SharePoint Services collaboration environment:

  http://technet.microsoft.com/en-us/library/cc261795.aspx

EMC Publications

- EMC Solutions for Microsoft Office SharePoint Server 2007 on VMware ESX Server EMC Celerra NS Series over NFS:

  http://www.emc.com/collateral/hardware/technical-documentation/h4381-emc-sol-sharepoint-ns-series-nfs.pdf

- Reference Architecture: EMC Solutions for Microsoft Office SharePoint Server EMC Celerra Unified Storage Platforms:

  http://www.emc.com/collateral/hardware/technical-documentation/h4150-emc-sol-ms-office-sharepoint-svr-2007-emc-celerra-na-iscsi-ra.pdf

- Reference Architecture: EMC Virtual Architecture for Microsoft SharePoint Server 2007 Enabled by EMC CLARiiON CX3-40, VMware ESX Server 3.5 and Microsoft SQL Server 2005:

  http://www.emc.com/collateral/hardware/technical-documentation/h4456-emc-vrtl-arch-ms-sharepoint-srvr-2007-ref-arch.pdf

VMware, Inc. 3401 Hillview Ave Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com