HITACHI
Inspire the Next

Microsoft
GOLD CERTIFIED
Partner

Virtualizing Microsoft® Office SharePoint® Server 2007 with VMware vSphere 4 on the Hitachi Adaptable Modular Storage 2000 Family

Reference Architecture Guide

By Steven Burns and Ken Ewers

February 2010

Hitachi Storage Solutions

Hitachi Data Systems

Summary

Many companies are turning to Microsoft® Office SharePoint® Server 2007 to manage content and create collaborative work environments. But deploying SharePoint can be a costly and complex endeavor without proper planning and without a storage solution and computing infrastructure that facilitates centralizing data and scaling the infrastructure as content and users increase.

The Hitachi Adaptable Modular Storage 2000 family provides a reliable, flexible, scalable and cost-effective modular storage system for SharePoint deployments. The 2000 family is ideal for more demanding application requirements and delivers enterprise-class performance, capacity and functionality at a midrange price. By deploying SharePoint on the 2000 family in a VMware environment, IT organizations gain additional advantages, such as increased server utilization resulting in lower costs and increased ease of management.

This paper provides best practices required to successfully deploy SharePoint 2007 in a server virtualized environment using VMware vSphere 4 configurations on the Hitachi Adaptable Modular Storage 2000 family. The solution is intended for use by IT administrators responsible for Microsoft SharePoint 2007, Microsoft SQL Server 2008, server virtualization and storage.

For best results use Acrobat Reader 8.0.

Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an e-mail message to SolutionLab@hds.com. Be sure to include the title of this white paper in your e-mail message.

Table of Contents

Virtualizing Microsoft® Office SharePoint® Server 2007 with VMware vSphere 4 on the Hitachi Adaptable Modular Storage 2000 Family

Reference Architecture Guide

Faced with explosive growth in unstructured data that needs to be shared among many employees, departments and divisions that are often operating in different time zones, countries, and continents, companies are turning to Microsoft® Office SharePoint® Server 2007 to manage content and create collaborative work environments. But deploying SharePoint can be a costly and complex endeavor without proper planning and without a storage solution and computing infrastructure that facilitates centralizing data and scaling the infrastructure as content and users increase.

The Hitachi Adaptable Modular Storage 2000 family provides a reliable, flexible, scalable and cost-effective modular storage system for SharePoint deployments. The 2000 family is ideal for more demanding application requirements and delivers enterprise-class performance, capacity and functionality at a midrange price. By deploying SharePoint on the 2000 family in a VMware environment, IT organizations gain additional advantages, such as increased server utilization resulting in lower costs and increased ease of management.

This paper provides best practices required to successfully deploy SharePoint 2007 in a server virtualized environment using VMware vSphere 4 configurations on the Hitachi Adaptable Modular Storage 2000 family. The solution is intended for use by IT administrators responsible for Microsoft SharePoint 2007, Microsoft SQL Server 2008, server virtualization and storage. It assumes familiarity with vSphere 4 and vCenter Server, Hitachi Storage Navigator Modular 2 software, Windows 2008, SharePoint 2007, and SQL Server 2008. Although the documented solution in this paper used the Hitachi Adaptable Modular Storage 2100 for testing, any member of the Hitachi Adaptable Modular Storage 2000 family provides a reliable, flexible and cost-effective storage platform for supporting demanding applications like SharePoint 2007.

Solution Components

This section describes the key components of this solution.

Hitachi Adaptable Modular Storage 2000 Family

The Hitachi Adaptable Modular Storage 2000 family is the only midrange storage product with symmetric active-active controllers that provide integrated, automated hardware-based front-to-back-end I/O load balancing. Both controllers in a 2000 family storage system are able to dynamically and automatically assign the access paths from the back of the controller to the LU. All LUs are accessible regardless of the physical storage front-end port or the server from which the access is requested. Utilization rates of each controller are monitored so that a more even distribution of workload between the two controllers can be maintained. When coupled with VMware round-robin load balancing, the 2000 family eliminates many complex and time consuming path planning tasks that storage administrators typically face.

No other midrange storage product that scales beyond 100TB has a serial attached SCSI (SAS) drive interface. The point-to-point back end design virtually eliminates I/O transfer delays and contention associated with Fibre Channel arbitration and provides significantly higher bandwidth and I/O concurrency.

Although the Hitachi Adaptable Modular Storage 2100 was used for this solution, the information in this paper is relevant for the other 2000 family members with the proper changes to account for capacity and performance differences. The 2000 family is an easy-to-use, scalable, cost effective storage system for mission-critical business applications like SharePoint Server 2007 and SQL Server 2008. It is also a top choice for tiered and standalone storage, consolidation, business continuity, data replication, backup and archiving. The Adaptable Modular Storage 2100 offers a rich set of features in a model that scales to 159 disk drives and delivers enterprise-class performance and capabilities at a modular price. Table 1 lists some of the 2100's specification options. For more information about the other member of the 2000 family, see the Hitachi Adaptable Modular Storage 2000 family Web site.

Table 1. Hitachi Adaptable Modular Storage 2100 Specification Options

Raw Capacity	156TB SATA 52TB SAS
Internal Disk Drives (SAS unless otherwise noted)	300GB (15K RPM) 450GB (15K RPM) 1TB SATA II (7200 RPM)
Maximum number of drives	159
Disk Drive Interfaces	SAS and SATA
Host Interfaces	Fibre Channel: 4Gb/sec iSCSI: GigE
Maximum Host Connections	4 Fibre Channel or 4 iSCSI
Maximum Attached Hosts Through Virtual Ports	512
SAS Links	16
Maximum Number of LUs	2,048
Maximum LU Size	60TB
Controller Cache (per system)	4GB to 8GB

Hitachi Dynamic Provisioning Software

On Hitachi Adaptable Modular Storage 2000 family systems, Hitachi Dynamic Provisioning software provides a dynamic provisioning feature with thin provisioning and wide striping components that provide virtual storage capacity to eliminate application service interruptions, reduce costs and simplify administration, as follows:

- Optimizes or "right-sizes" storage performance and capacity based on business or application requirements.
- Supports deferring storage capacity upgrades to align with actual business usage.
- Simplifies and adds agility to the storage administration process.
- Provides performance improvements through automatic optimized wide striping of data across all available disks in a storage pool.

For more information, see the Hitachi Dynamic Provisioning Software datasheet.

The wide-striping technology that is fundamental to Hitachi Data Provisioning software dramatically improves performance, capacity utilization and management of your environment. Storage is allocated from a Hitachi Dynamic Provisioning pool as V-VOLs. By deploying SharePoint 2007 and SQL Server 2008 using V-VOLs from Hitachi Dynamic Provisioning storage pools on the 2000 family, you can expect the following benefits:

- An improved I/O buffer to burst into during peak usage times or intense maintenance activities like content indexing or database integrity checks
- A smoothing effect to the SQL Server workload that can eliminate hot spots across the different RAID groups, resulting in reduce data migration moves related to performance or capacity constraints
- Elimination of excess, unutilized capacity by leveraging the combined capabilities all disks comprising a storage pool
- Elimination of the need to manage the placement of heavy load databases

VMware vSphere 4

VMware vSphere 4 can help reduce hardware footprints and capital expenses dramatically through server consolidation. Utilizing VMware products and features such as ESX, vCenter Server, High Availability (HA), Distributed Resource Scheduler (DRS) and Fault Tolerance (FT), vSphere allows for a robust environment, centralized management and gives administrators control over key capabilities.

VMware provides product features that can help manage the entire infrastructure:

- **VMotion** — Allows for non-disruptive migration of both virtual machines and storage. Its performance graphs allow for monitoring of resources, virtual machines, resource pools and server utilization.
- **Distributed Resource Scheduler (DRS)** — Monitors resource utilization and intelligently allocates system resources as needed.
- **High Availability (HA)** — Monitors for hardware and OS failures and automatically restarts the virtual machine, providing cost effective failovers.
- **Fault Tolerance (FT)** — Provides continuous availability for applications by creating a live shadow instance of the virtual machine that stays in lockstep with the primary instance. If there is a hardware failure the shadow instance instantly takes over and eliminates even the smallest data loss.

Microsoft Office SharePoint Server 2007

SharePoint Server 2007 is an integrated collaboration application that allows organizations to share and manage content using intranet and extranet portals and meeting workspaces that are easy to create and administer. SharePoint can be integrated with a wide variety of Microsoft applications including Word, Project Server, Excel, and others.

Microsoft SQL Server 2008

Microsoft SQL Server 2008 facilitates the management of any data, any place and any time. Together with the 2000 family, SQL Server 2008 provides a scalable, high-performance database engine for any midrange to enterprise level applications that require the highest levels of availability and security, while reducing the total cost of ownership through enhanced enterprise-class manageability.

Storage Navigator Modular 2 Software

Hitachi Storage Navigator Modular 2 software enables essential functions for the management and optimization of individual Hitachi storage systems. It provides both a Web-accessible graphical management interface and a command line interface (CLI) to allow ease of storage management.

Use Storage Navigator Modular 2 software to map security levels for SAN ports and virtual ports and for inter-system path mapping, for RAID-level configurations, for logical unit (LU) creation and expansion and for online Volume Migrations.

Hitachi Performance Monitor Feature

Hitachi Performance Monitor feature is part of the Storage Navigator Modular 2 software package. It acquires information on the performance of RAID groups, logical units (LUs) and other elements of the storage system while tracking the utilization rates of resources such as hard disk drives and processors. Information is displayed with line graphs in the Performance Monitor windows or saved in `.csv` files that you can analyze later. Use Performance Monitor to validate that the underlying storage system is not taxed by the SQL Server environment.

Hitachi Device Manager Software

Hitachi Device Manager software provides centralized management of all Hitachi storage systems, including the Universal Storage Platform family. Device Manager software can link to Storage Navigator software and has the ability to provision using storage pools, manage replication between storage systems and logically group resources for more efficient management. It also provides the device interface into Hitachi storage systems for the heterogeneous Hitachi Storage Services Manager software and other third-party storage management tools. Device Manager software's Provisioning Assistant enables you to integrate and manage various models and types of storage systems as a single, virtual storage pool, allowing you to do more with less.

Hitachi Tuning Manager Software

Hitachi Tuning Manager software enables you to proactively monitor, manage and plan the performance and capacity for the Hitachi modular storage that is attached to your vSphere servers. Hitachi Tuning Manager software consolidates statistical performance data from the entire storage path. It collects performance and capacity data from the operating system, SQL Server instances, switch ports, and storage system information such as RAID groups and LUs IOPS and latency measures. Hitachi Tuning Manager software provides the administrator a complete performance picture of the environment. It provides historical, current and forecast views of these metrics. For more information about Hitachi Tuning Manager software, see the Hitachi Data Systems support portal.

Reference Architecture

This reference architecture is designed to support a SharePoint 2007 farm with 1TB in total user content stored on SQL Server 2008. Ten 100GB SharePoint content databases are used for the user data. The reference architecture uses four Web front end (WFE) servers to handle user access and queries, three application servers to handle all of SharePoint roles including the index and Excel services roles, and one server with SQL 2008 Server installed to handle the database requirements of the environment. Figure 1 shows the physical layout of the reference architecture environment.

Figure 1. Physical Layout of the Reference Architecture

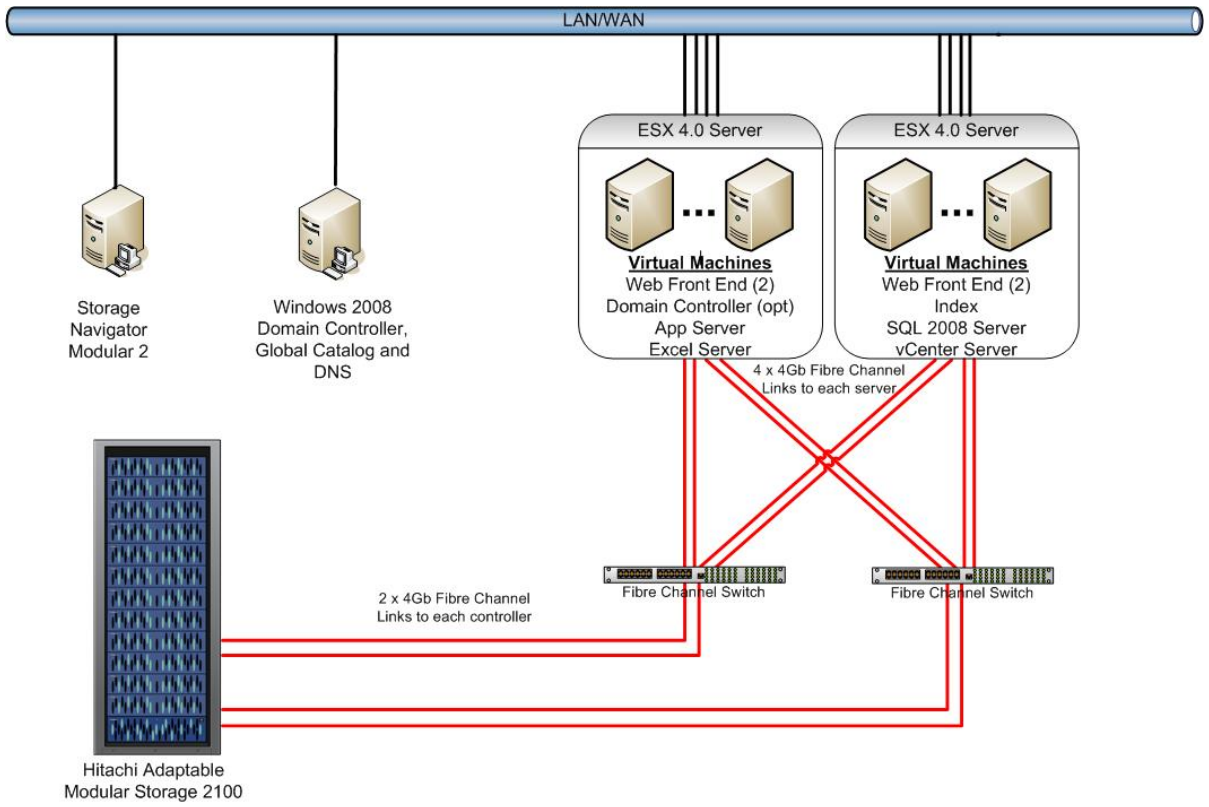

This reference architecture includes two VMware ESX 4.0 hosts that house a total of ten virtual machines. All of the servers run Windows 2008 Service Pack 2. Seven of the virtual machines have SharePoint Server 2007 installed, one has SQL Server 2008 Service Pack 1, one acts as an Active Directory domain controller DC for the environment, and one hosts VMware vCenter, which is used to manage the vSphere 4 environment.

Software

Table 2 lists the software components of this solution.

Table 2. Software Resources

Software	*Version*
VMware ESX	4.0.0
VMware vCenter	4.0
Windows Server	2008 Enterprise edition, Service Pack 2
SQL Server	2008 Enterprise edition, Service Pack 1
SharePoint Server	2007 Enterprise edition, Service Pack 2
Hitachi Storage Navigator Modular 2	8.3*

*Note: To use Hitachi Dynamic Provisioning software in this solution, the license for the Hitachi Dynamic Provisioning software is also required. The license can be installed using Hitachi Storage Navigator Modular 2 software's GUI or command-line interface.

Hardware

Table 3 lists the hardware resources for this solution.

Table 3. Hardware Resources

Hardware	*Configuration*
Hitachi Adaptable Modular Storage 2100	Microcode 0883A* 2 controllers 2 4Gb Fibre Channel ports per controller 2 disk trays 23 450GB 15K RPM SAS disks 2GB cache per controller
Brocade 48000 FC Director	12 4Gb Fibre Channel ports FOS 5.3.1a
Two Dell R905 servers	BIOS firmware 4.0.3** 4 quad core AMD Opteron 1.9GHz processors per server 64GB of memory per server 2 Emulex LPe11002 4Gb Dual Port Fibre Channel HBAs per server 4 GigE network interface ports per server

*Use the latest or the most recent Hitachi supported levels of microcode and FOS levels. Specific minimum levels of microcode are required to support certain features.
**Use the latest available or VMware supported BIOS firmware level.

An additional server was used in the lab environment to host management applications such as Hitachi Storage Navigator Modular 2 and the vSphere client. This server was also used to remote desktop into the virtual machines for application installation and configuration tasks.

In the lab environment that this reference architecture is based on, an Active Directory domain controller was running on a virtual machine. In a production environment a domain controller running as a virtual machine might not be required if one or more already exist.

Storage Architecture

This reference architecture uses a storage design that consists of two Hitachi Dynamic Provisioning pools for the SQL and SharePoint data. Each Dynamic Provisioning pool consists of two RAID-1+0 (2D+2D) groups using 450GB 15K drives. One RAID-5 (4D+1P) group handles the virtual machine files and boot disks. When designing the Dynamic Provisioning configuration for a SharePoint 2007 environment, Hitachi Data Systems recommends using two storage pools. Use one pool for the SQL databases required for SharePoint (content databases, configuration databases, search databases and others) and use the other pool for database logs other content such as the flat files for the index and query servers. Using two pools instead of a single combined pool ensures that the environment meets Microsoft's recommendation to have a database and its log files on separate sets of disk spindles.

The environment was built using Hitachi Dynamic Provisioning software. Two pools of two RAID-1+0 (2D+2D) groups were created for the environment. Dedicated volumes within the pool simplify backup operations due to their granularity and also provide an easier way to monitor the capacity utilization of each of the volume types. Hitachi Data Systems recommends using a RAID-1+0 (2D+2D) configuration as the building block for your Dynamic Provisioning pool as it provides good reliability, performance, and good capacity increment for your pool. While it is possible to create a pool with other RAID group types supported by the 2000 family, RAID-1+0 (2D+2D) provides the required protection level that Microsoft recommends for a SQL Server environment.

Figure 3 shows the physical layout of the Hitachi Adaptable Modular Storage 2100 used in this solution.

Figure 3. Storage Design for Solution

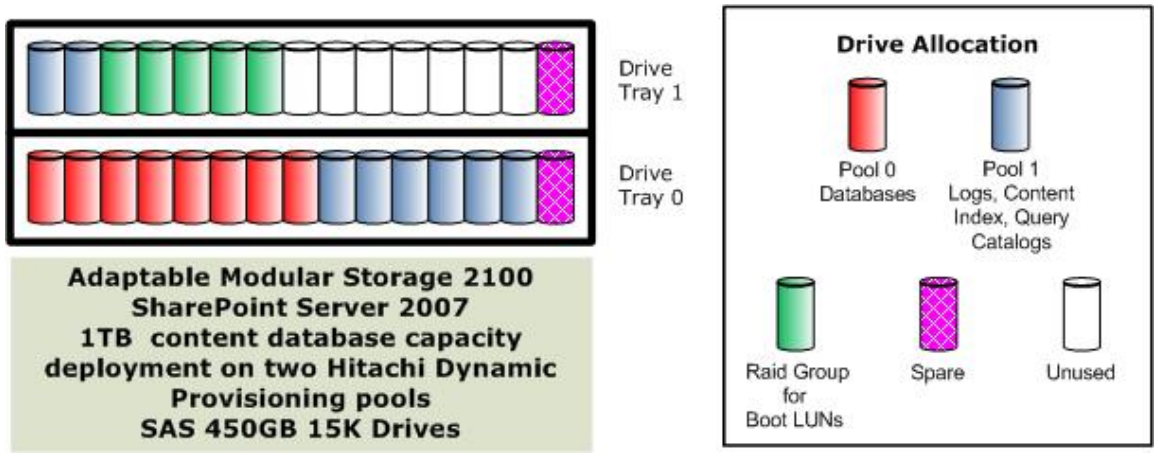

When implementing a solution in a vSphere4 environment that uses Hitachi Dynamic Provisioning software on the storage, it is important to understand the relationship between the following components:

- **NTFS volumes** — From a Windows 2008 perspective, a partition on a disk drive. In a Dynamic Provisioning environment, it is important to note that if you do a full format on a NTFS volume all of the space is allocated on the 2100. For this reason, Hitachi Data Systems recommends performing quick formats on NTFS volumes to take advantage of Hitachi Dynamic Provisioning software's thin provisioning.
- **Windows 2008 Disk Drive** — The disk drives on a virtual machine in a VMware environment look and act like a local physical hard drive from a Windows 2008 perspective. Utilities such as DiskPart work the same. If you look at the properties of a disk drive on a VMware virtual machine from Disk Management, they are shown as a VMware virtual disk SCSI disk device. From a VMware perspective, a virtual machine's virtual disk exists as a `.vmdk` file on a VMware datastore.
- **VMware `.vmdk` files** — These are the files created and stored on a VMware datastore where a virtual machine's hard drive is located. The size of these files can be increased as needed by editing the settings for the virtual machine. A `.vmdk` file can have one of three formats: thin, thick, or eagerzeroedthick. For more information, see the "VMware Configuration" section in this paper. Thick (also known as zeroedthick) is

recommended for Dynamic Provisioning environments because, unlike the eagerzeroedthick format, it does not write to all of the allocated space when the `.vmdk` is created.

- **VMware datastores** — A managed object in a VMware data center. It represents a storage location for virtual machine files. A datastore can be a VMFS volume, a directory on a NAS device, or a local file system path. For this solution, a datastore is a VMFS volume on SAN storage presented as an LU from the 2100. Because the scope of a datastore is at the data center level, all hosts within a VMware data center see the same datastores. This solution uses a one-to-one relationship between LUs on the Hitachi Adaptable Modular Storage 2100 and VMware datastores.
- **LUs created on the 2100** — For this solution, LUs are allocated from an Dynamic Provisioning pool on the 2100 and presented to both hosts to be used for VMware datastores. LUs must be presented to both hosts so that virtual machines can be moved between ESX hosts as necessary.
- **Hitachi Dynamic Provisioning pools** — Dynamic Provisioning pools are a licensed feature of the 2100 and offer the benefits of thin provisioning and wide striping. For a SharePoint 2007 environment, use a minimum of two Dynamic Provisioning pools. This allows the SQL databases and their related log files to be stored on separate sets of spindles in accordance with Microsoft recommendations. It is important to note that scenarios exist within a SharePoint 2007 virtualized on vSphere4 environment where thin provisioning benefits are defeated, including these:
  - **Full format of an NTFS volume** — Because an NTFS full format writes to all blocks on the volume the total space of the volume is allocated from the Dynamic Provisioning pool.
  - **Use of eagerzeroedthick format on a `.vmdk` file** — This format prezeroes all of the allocated space. This format is required for certain VMware features such as fault tolerance.

LUs from the 2100 are allocated to the ESX hosts and formatted as virtual machine file system (VMFS) volumes from which virtual disks are created. The virtual disks (vDisks) are presented to the Windows Server 2008 guest operating system and can be partitioned and used in NTFS file systems for the SQL Server 2008 databases, `tempdb`, and transaction logs required for the SharePoint environment and for the SharePoint index and query files. This solution is designed to support 10 SharePoint content databases of 100GB each for a total capacity of 1TB database. While Microsoft recommends limiting content databases for SharePoint 2007 to 100GB in size, this is not a hard limit.

When planning the storage configuration for SharePoint running in a virtual environment, keep these kinds of data in mind:

- **SQL database files** — These include the databases for the SharePoint content databases, the search databases, SSP database, configuration database and others. Place all of the SQL databases in one Dynamic Provisioning pool.
- **SQL log files** — Each database has a related log file. Store these files in the second Dynamic Provisioning pool to separate the logs onto a different set of disks from their related databases. Sizing is dependent on activity on the database and backup procedures. Size the SQL log volumes based on either calculated or measured values for the amount of logs generated per day taking into account how many days' logs might need to be stored. This is dependent both on backup procedures and schedules as well as the historical reliability of the backup process for the environment.
- **SQL `tempdb`** — Place this file on the same Dynamic Provisioning pool as the logs. This is a critical file and proper configuration is important.
- **Content index flat files (for SharePoint servers running the index or query services)** — These files are not SQL related and reside on servers running the Index or Query services in a SharePoint 2007 farm. Storage requirements for these are discussed further in the Sizing section below.
- **Boot LUs for virtual machines** — The storage required for each virtual machine varies depending upon the applications installed. VMware recommends allocating 16GB to 32GB per virtual machine. The size of the disk can be increased at any point by editing the virtual machine's settings. This solution stores these files on a dedicated RAID-5 (4D+1P) group.

Sizing LUs and Dynamic Provisioning Pools

- **SharePoint content databases** — This solution uses ten 100GB content databases. The actual number and size of content databases in your environment vary depending upon the amount of content and the structure and hierarchy of your environment. Microsoft recommends that content databases be limited to 100GB in size but this is not a hard limit and many production SharePoint environments exist with larger content databases. Microsoft also recommends that an additional 20 percent of capacity be allocated to avoid out-of-space issues caused by unexpected growth. Based on these recommendations, base sizing calculations on 84GB per content database. If additional space is needed, Microsoft recommends adding SharePoint site collections and content databases.
- **SharePoint Search database** — The size of the search database can vary greatly depending on the number and average size of files that are indexed. Microsoft's calculation for the size of the search database follows:

  GB of disk space required = Total_Content_Size (in GB) x File_Size_Modifier x 4

  File_Size_Modifier is one of the following based on the average size of the files stored in the content databases, as follows:

  – 1.0 if average files size of the content is 1KB

  – 0.12 if average file size of the content is 10KB

  – 0.05 if the average file size of the content is 100KB or larger

  This calculation can result in range of from 200GB to 4TB for a total content size of 1TB depending on the average size of files in the content databases.
- **Other SharePoint databases** — Many other databases are created when you install and configure SharePoint. These can usually be stored in a single .vmdk file.
- **Database log files** — **Estimate t**he storage capacity required for the logs at 20 percent of the size of the databases. Backup procedures need to be taken into account as well as historical reliability of the backups.
- **SQL `tempdb`** — Estimate the storage capacity required for `tempdb` files at 10 percent of the total size of the databases.
- **Content index files** — The content index is a flat file stored on the index server. Use the following formula to calculate the storage capacity needed for this file:

  GB of disk space required = Total_Content_Size (in GB) x File_Size_Modifier x 2.85

  File_Size_Modifier is one of the following based on the average size of the files stored in the content databases.

  – 1.0 if average files size of the content is 1KB

  – 0.12 if average file size of the content is 10KB

  – 0.05 if the average file size of the content is 100KB or larger
- **Query catalog files** — A copy of the query catalog file is stored on each of the servers hosting the query role. The query catalog file is a flat file that is a copy of the content index that is propagated to each of the query servers. This file requires approximately the storage capacity that the content index file requires.
- **Boot LUs** — The boot LUs for the virtual machines in this reference architecture are configured as 40GB vDisks except for the SQL 2008 server which uses an 80GB vdisk.

LU Layout

Table 4 shows the LU layout for the reference architecture.

Table 4. LU Layout

Dynamic Provisioning Pool	LU Number	Purpose	Capacity	Server Assignment
Pool 000	0000	SharePoint search DB	500GB	SQL
Pool 000	0001	SharePoint content DB 1	100GB	SQL
Pool 000	0002	SharePoint content DB 2	100GB	SQL
Pool 000	0003	SharePoint content DB 3	100GB	SQL
Pool 000	0004	SharePoint content DB 4	100GB	SQL
Pool 000	0005	SharePoint content DB 5	100GB	SQL
Pool 000	0006	SharePoint content DB 6	100GB	SQL
Pool 000	0007	SharePoint content DB 7	100GB	SQL
Pool 000	0008	SharePoint content DB 8	100GB	SQL
Pool 000	0009	SharePoint content DB 9	100GB	SQL
Pool 000	0010	SharePoint content DB 10	100GB	SQL
Pool 000	0011	Other SharePoint DBs*	150GB	SQL
Pool 000	0012	`tempdb`	125GB	SQL
Pool 001	0013	Logs**	250GB	SQL
Pool 001	0014	SharePoint content index	300GB	Index
Pool 001	0015	WFE1 query catalog	150GB	WFE1
Pool 001	0016	WFE2 query catalog	150GB	WFE2
Pool 001	0017	WFE3 query catalog	150GB	WFE3
Pool 001	0018	WFE4 query catalog	150GB	WFE4

*Note: LU 11 is used for all SharePoint databases other than the content databases, SharePoint search database, and `tempdb`. This includes the SharePoint configuration database, the SharePoint admin content database, the SSP admin (Site Collection) Content database, the SSP database and the WSS search database. Some environments might include other databases.

**Note: LU 13 is used to store all of the database transaction logs.

Scaling the Storage Configuration

If an environment requires larger amounts of storage either initially or due to growth, additional capacity can easily be added to the existing Dynamic Provisioning pools. Capacity is added to an existing pool by an increment that is dependent on the drive size and underlying RAID structure of the pool. For the configuration documented in this solution, the increment is approximately 900GB, based on 450GB drives in a RAID-1+0 (2D+2D) configuration.

Other RAID Types and SATA Disks

While this paper documents a storage layout using SAS drives in a RAID-1+0 (2D+2D) configuration, in some scenarios other types of RAID configuration or the use of SATA drives might be feasible. As done in this solution in most cases RAID-5 can be used for the storage allocated for the datastores housing the virtual machine OS vDisks. According to Microsoft you might also be able to use RAID-5 for databases with static Web content where random read performance is important; however, use RAID-1+0 for any database used for collaboration.

Depending on your specific environment, you might use more than two Dynamic Provisioning pools, for example, in these circumstances:

- Use of RAID-5 for specific purposes.
- I/O requirements necessitate the use of additional pools.
- Use of multiple disk sizes.
- Sharing 2000 family storage with other applications.

SharePoint Capacity Planning Tool

Use the SharePoint Capacity Planning Tool to help plan your storage configuration. Use the tool can be used to model both SharePoint and WSS environments. To obtain and use this tool, download System Center Capacity Planner 2007 (SCCP 2007) first and then download the SharePoint Capacity Tool. This tool is designed for standalone server environments but can also help with planning a virtualized SharePoint 2007 farm.

For more information, see the Microsoft TechNet article "SharePoint Capacity Planning Tool Executive Overview."

vSphere 4 Architecture

The vSphere 4 environment for this reference architecture consists of two ESX 4.0 servers hosting a total of 10 virtual machines.

For more information, see *Mastering VMware vSphere 4* by Scott Lowe, ISBN:978-0-470-48138-7.

Sizing ESX Hosts

Before you can establish the specifications for your ESX hosts, you must first determine how many virtual machines will run on each ESX host. You must also determine the virtual CPU, memory and capacity requirements for the virtual machines and consider whether the ESX host will be part of a cluster and whether features like VMotion will be used. If you are sizing your environment for failover or HA scenarios, verify that each host has enough resources to support all of the virtual machines in the environment.

Virtual CPUs

A virtual environment contains both virtual CPUs (vCPUs) and physical CPUs. Virtual CPUs are the processors within the virtual machines, and they equate to one physical processor core. The virtual machine monitor is responsible for virtualizing the CPUs. VMware recommends ensuring that the total number of vCPUs assigned to all the virtual machines is equal to or less than the total number of cores on the physical ESX host. Keep in mind that the ESX service console requires one CPU and at least 400MB of memory.

vSphere Memory Configuration Guidelines

vSphere manages memory at the hypervisor level, enabling such features as memory sharing across VMs that have the same guest operating systems, memory over-commitment and a memory balloon technique that allows virtual machines that are not using all of their allocated memory to give memory to virtual machines that require more.

When sizing an ESX host for memory, consider two additional factors beyond what is needed for the virtual machines:

- **ESX host memory overhead —** Overhead for the service console, the VMkernel and each virtual machine
- **Virtual machine memory settings —** Overhead for the ESX host, also called guest memory, which can be broken down into the following parts:
  - **Configured memory** — Memory size assigned to the virtual machine
  - **Touched memory** — Amount of virtually allocated memory the virtual machine is using
  - **Reserved memory** — Portion of configured memory vSphere reserves for the virtual machine, taking into account overhead
  - **Swappable memory** — Portion of configured memory beyond reserved, which can be reclaimed by the balloon driver or by vSphere swapping

Figure 4 illustrates virtual machine memory settings.

Figure 4. Virtual Machine Memory Settings

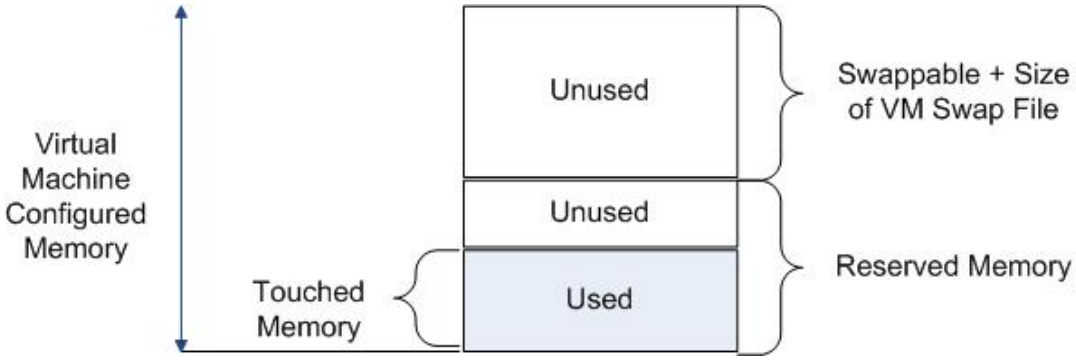

In a vSphere4 environment the LUs from the 2100 are presented to the virtual machines running on an ESX host as either a vDisk created as a .vmdk file on a datastore or as a mapped raw LU.

VMware Disk Management

ESX hosts can access LUs in two ways, through a VMware virtual machine file system (VMFS) datastore or via a raw device mapping (RDM). Virtual disks are files stored on a VMFS datastore, which is a logical container, in the form of `.vmdk` files. The VMFS datastores are deployed on storage devices using the VMFS format. VMFS is optimized for VMs and hides the specifics of the underlying storage.

VMFS datastores can be accessed by several ESX hosts and the cluster feature allows for distributed file locking for the VMs. The VMFS datastore can be extended while the client is running and can extend across multiple LUs. Using Storage VMotion, the `.vmdk` file can be moved to another VMFS datastore nondisruptively.

As an alternative, LUs can be mapped as a raw device mapping (RDM). This feature allows a LU to be mapped directly to a VM. RDMs are useful for command devices and any other device that requires direct communication to the storage processor. Hitachi Data Systems recommends against using RDM for SQL Server environments because the enhanced performance and features of the VMFS are not available. However, keep in mind that when using Microsoft cluster services (MSCS), you must use RDM in physical compatibility mode.

For more information about VMFS and RDM, see VMware's *Fibre Channel SAN Configuration Guide*.

It is important to understand the various types of disk devices that might exist in a VMware environment:

- **Virtual disks** — These are virtual disks created on a VMFS datastore and stored as `.vmdk` files. These are also referred to as vDisks. Virtual disks can come in the following formats:
  - **Thin** — Allocates space as needed so the size of the `.vmdk` file on the datastore is only what is actually used. Hitachi Data Systems recommends against using thin format when using Hitachi Dynamic Provisioning software on the 2100 because thin provisioning is handled on the storage system.
  - **Thick (zeroedthick)** — Creates a vmdk file on the datastore that is the size of the virtual disk created but does not "prezero" the space. Using this format if you create a virtual disk of 300GB and place 100GB of data in it, the vmdk file on the datastore will be 300GB but the space actually used on the 2100 will be 100GB. This is the format recommended when using Hitachi Dynamic Provisioning software on the 2100.
  - **Eagerzeroedthick** — "Prezeroes" all of the space allocated. With this format, a 500GB virtual disk has a 500GB `.vmdk` file on the datastore and it also uses 500GB of space on the 2100, even when using Hitachi Dynamic Provisioning software. While this format is not "thin-friendly," it is required in certain circumstances. For example, if VMware fault tolerance is enabled on a thin or thick vDisk, they are converted to Eagerzeroedthick because this format is required for fault tolerance on VMFS.
- **Mapped Raw LUNs** — Raw device mapping (RDM) storage devices. These can be configured in two different modes, physical compatibility mode (pRDM) or virtual mode (vRDM). A detailed explanation of RDMs is outside the scope of this paper. It is important to note that RDMs are required instead of vDisks for the SQL server in a SharePoint farm environment if it is clustered using Microsoft clustering.

Application Architecture

This solution's SharePoint 2007 farm environment architecture consists of eight virtual machines on two VMware ESX 4.0 hosts. Seven of these VMs have SharePoint Server 2007 installed and the eighth VM has SQL Server 2008 SP1 installed. A SharePoint environment always consists of three tiers, as follows:

- Web
- Application
- Database

Any given server in the environment can support one or more of these tiers. The environment tested for this solution uses four Web front end (WFE) servers that are configured with the Windows network load balancing service (NLBS). These four WFEs also have the query service enabled and hold a copy of the query catalog. The application tier consists of the application server, the Excel services server, and the index server. The database tier consists of the eighth server, which hosts SQL Server 2008 SP1.

The virtual machines and the ESX hosts that host them are shown in Table 4.

Table 4. ESX Hosts and Virtual Machines

ESX Host	*Virtual Machine*	*Resources*	*Operating System*	*Applications and Services*
ESX1	WFE1	2 vCPU 4GB memory	Windows 2008 Service Pack 2 64-bit	SharePoint 2007 Service Pack 2 Web application service Search query service Windows NLBS
ESX1	WFE2	2 vCPU 4GB memory	Windows 2008 Service Pack 2 64-bit	SharePoint 2007 Service Pack 2 Web application service Search query service Windows NLBS
ESX1	Application server	2 vCPU 4GB memory	Windows 2008 Service Pack 2 64-bit	SharePoint 2007 Service Pack 2 Web application service
ESX1	Excel server	2 vCPU 4GB memory	Windows 2008 Service Pack 2 64-bit	SharePoint 2007 Service Pack 2 Web application service Excel calc service
ESX2	WFE3	2 vCPU 4GB memory	Windows 2008 Service Pack 2 64-bit	SharePoint 2007 Service Pack 2 Web application service Search query service Windows NLBS
ESX2	WFE4	2 vCPU 4GB memory	Windows 2008 Service Pack 2 64-bit	SharePoint 2007 Service Pack 2 Web application service Search query service Windows NLBS
ESX2	SQL Server	4 vCPU 16GB memory	Windows 2008 Service Pack 2 64-bit	SQL Server 2008 Service Pack 1
ESX2	Index Server	4 vCPU 8GB memory	Windows 2008 Service Pack 2 64-bit	SharePoint 2007 Service Pack 2 Web Application Service Search Indexing Service

The lab environment also included two additional virtual machines. The Windows Domain Controller for the environment was hosted on ESX1 and the vCenter server was hosted on ESX2. These and other infrastructure and management servers might already exist in a production environment and additional instances might not be needed to implement this solution.

Sizing Virtual Machines

To minimize time required to build your SharePoint environment, use templates to create virtual machines whenever possible. All of the VMs in the solution environment are Windows 2008 Server 64 bit with the exception of the VM that hosts VMware vCenter.

Create the VM template with two vCPUs and 2GB of memory. Install Windows 2008 Server 64 bit on the template and install all Windows roles and features required for SharePoint 2007. Allocate 40GB for the vDisk for the boot LU.

When creating the SharePoint VMs from the template, follow these initial sizing guidelines:

- **WFE (Web Front End) VMs –** Create these from the VM template. Start with two vCPUs and 4GB of memory and scale up as necessary. Because the WFE and query roles are CPU intensive, additional vCPUs might be required. If the WFEs also run the query service, create a separate vDisk to hold the query catalog, which is a copy of the content index that is propagated from the Index server. Attach this vDisk via a separate vSCSI adapter.
- **Index server —** Create this server from the template. Start with two vCPU and 6GB of memory and scale up as necessary. Allocate a second vDisk attached to a second vSCSI adapter for the content index. This vDisk needs to be greater than twice the size of the actual content index due to the process the indexing service uses to merge data. The size of the content index is also dependent on the average size of the files being index. For smaller files (more files per GB of content), the content index is larger and conversely, for larger files (fewer files per GB of content) the content index is smaller. Microsoft recommends multiplying the estimated size of the content index by 2.85.
- **Application server —** This includes servers running the Excel service. Create these from the VM template. Allocate two vCPUs and 2GB of memory and scale up if necessary. These roles do not usually require any additional resources.
- **SQL server —** Create this server from the VM template. While VMware recommends allocating two vCPUs and 6GB of memory and scaling up as necessary additional memory increases the caching efficiency of the SQL server, which decreases the number of IOPs for storage and increases performance. For that reason, Hitachi Data Systems recommends allocating 12GB to 16GB of memory. If Microsoft clustering is used for the SQL server, see refer to the VMware's *Setup for Failover Clustering and Microsoft Cluster Service* guide. Note that you must use RDMs instead of vDisks for shared storage when using Microsoft clustering. Specific requirements depend on cluster design and OS version.
- **Memory and vCPU** — These requirements vary depending upon the specific environment.

Scaling the Application Architecture

When scaling the application architecture for this solution, or to modify it for high availability, consider the following:

- **WFE servers —** Web front end servers can either be scaled up or scaled out. To scale up, additional vCPUs, memory or both can be allocated to the virtual machine. Because both the WFE and query roles are CPU intensive, additional vCPUs are usually the better investment. To scale out additional WFE servers can be added but testing by Microsoft and VMware shows that diminishing returns are realized above a certain number of WFEs. Testing done by Microsoft on query server performance shows that adding query servers beyond a total of seven decreases the total capacity for queries per second. WFE servers are load balanced so that if one of them fails, users are directed to a server that is online.
- **Application servers —** These servers usually have relatively low resource requirements and can easily be scaled up or out to handle any additional load. Application servers can be added to a SharePoint environment as needed to fulfill various roles. Almost all roles enabled on an application server can be redundant and load balanced.
- **Excel server —** The Excel calculation service is enabled on an application server and usually does not need to be scaled up. Multiple application servers can have the Excel service enabled on them for high availability.

- **Index server —** The search index service is also enabled on an application server but you cannot have more than one index server for a Shared Service Provider (SSP). For more information about improving index server performance, see the Microsoft TechNet article "Change the Indexer Performance setting."
- **Database server —** Adding memory can improve performance of the SQL server by increasing the effectiveness of caching SharePoint data. For high availability in a production environment cluster the SQL server.

SAN Architecture

Figure 4 shows the Fibre Channel connections for the reference architecture.

Figure 4. Fibre Channel Connections

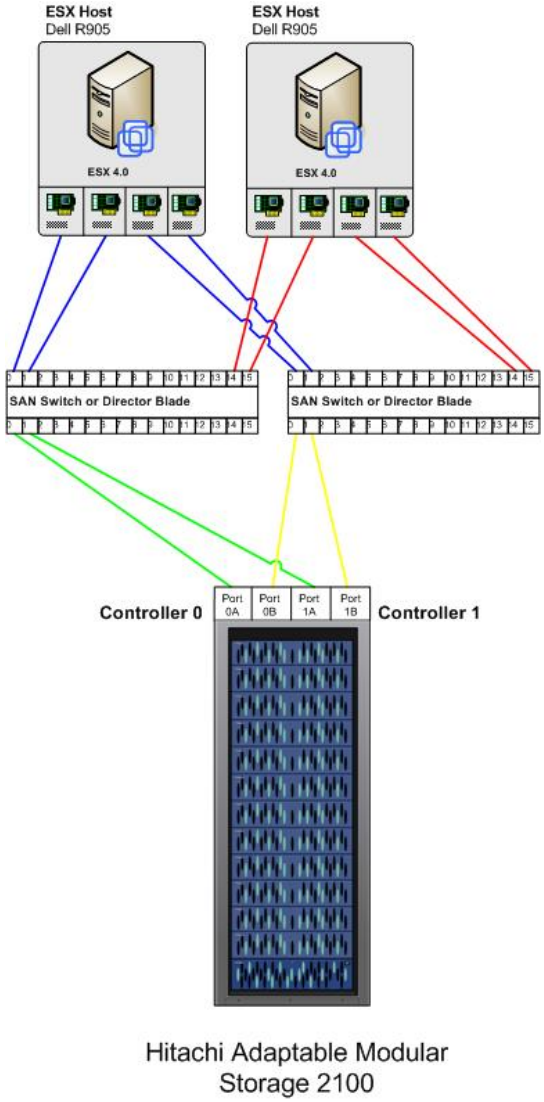

Table 5 lists the configuration of the paths between the HBA ports on the ESX host and the storage ports on the 2100. For the reference architecture, each row in the table is configured as a separate zone.

Table 5. Path Configuration for ESX Hosts

ESX Host	Host HBA Number	Director Zone Name	Storage System Port	Storage System Host Group
ESX 1	vmHBA 1	ESX_1_HBA1_AMS2100_0A	0A	ESX_1_vmHBA_1
ESX 1	vmHBA 2	ESX_1_HBA2_AMS2100_1A	1A	ESX_1_vmHBA_2
ESX 1	vmHBA 3	ESX_1_HBA3_AMS2100_0B	0B	ESX_1_vmHBA_3
ESX 1	vmHBA 4	ESX_1_HBA4_AMS2100_1B	1B	ESX_1_vmHBA_4
ESX 2	vmHBA 1	ESX_2_HBA1_AMS2100_0A	0A	ESX_2_vmHBA_1
ESX 2	vmHBA 2	ESX_2_HBA2_AMS2100_1A	1A	ESX_2_vmHBA_2
ESX 2	vmHBA 3	ESX_2_HBA3_AMS2100_0B	0B	ESX_2_vmHBA_3
ESX 2	vmHBA 4	ESX_2_HBA4_AMS2100_1B	1B	ESX_2_vmHBA_4

Multipathing on an ESX Host

To maintain a constant connection between the ESX hosts and storage, ESX supports multipathing. Multipathing allows multiple physical or logical connections from the host to the storage. To support multipathing, Hitachi Data Systems recommends that the physical host contain at least two HBAs that connect to at least one Fibre Channel port on each storage system controller. This reference architecture uses four HBA ports per ESX host.

In ESX, several types of multipathing policies are available through the VMware Native Multipathing Plugin (NMP):

- **Fixed (default) —** Uses the designated preferred path, if it is configured. Otherwise, it uses the first working path discovered at system boot time. If the host cannot use the preferred path, it selects a random alternative available path. The host automatically reverts to the preferred path as soon as that path becomes available. This is the default setting on ESX and requires manual load balancing and path distribution across SAN fabrics and storage controllers.
- **Round-robin (recommended) —** Uses a path selection algorithm that rotates through all available paths, enabling load distribution across the paths.

  Round-robin is the best choice for the Adaptable Modular Storage 2000 family of storage systems due to their symmetric active-active controller design. It ensures all resources within the storage system are used while maintaining the path failover capability of the environment and simplifying setup as all LUs are mapped to all ports assigned to an ESX host.
- **Most recently used (MRU) —** Selects the path the ESX host used most recently to access the given device. If this path becomes unavailable, the host switches to a different path and continues to use the new path while it is available.

Network Architecture

This reference architecture uses four network interface cards (NICs) on each ESX host. The number of NICs required for a specific environment varies based on multiple factors. The following sections describe some of the design considerations for both the physical and virtual networking.

Physical Network Configuration

The number of physical NICs required and supported in an ESX server is dependent on multiple factors, including these:

- **Limitations of the physical server —** Your physical server has a maximum number of network ports that can be supported. Interface speed and the network port configurations that VMware supports. are listed in the "Networking Maximums" section of the vSphere 4 *Configuration Maximums* guide.
- **vSphere features used —** Certain VMware features such as fault tolerance and vMotion require dedicated NICs.
- **Bandwidth required to support virtual machines —** Enough NICs are needed to support the planned number of virtual machines and the expected traffic.
- **Redundancy requirements —** Additional NICs might be required for redundancy purposes in some environments.
- **Use and configuration of VLANs —** The use of VLANs can reduce the number of NICs required for a vSphere 4 environment.

In a production environment, best practice is to use a minimum of six Giganet Ethernet NICs per ESX host but eight to twelve NICs per host is common especially in complex environments using physical servers with large amounts of CPU and memory resources. NICs are required for the following uses:

- **Service console —** One NIC is required, use two for redundancy.
- **Virtual machine traffic —** Allocate at least two NICs for VM traffic. Gigabit Ethernet or faster is recommended.
- **VMotion —** At least one NIC is required, use two for redundancy. Gigabit Ethernet or faster is recommended.
- **VMware Fault Tolerance —** If this feature is used it needs at least one NIC, two for redundancy. Gigabit Ethernet or faster is recommended.

VMware Network Configuration

Two types of virtual switches are available in a vSphere 4 environment. These are commonly referred to as vNetwork standard switches and vNetwork distributed switches. Within the vCenter GUI, these are referred to as virtual switches and distributed virtual switches (vDS). A standard switch is a virtual switch configured at the host level, and a vDS is configured at the vCenter level. A vDS operates as a single switch across multiple hosts and allows for easier migration of virtual machines between hosts. If standard switches are used, the switches on each host must be configured identically to support moving virtual machines between hosts. An in-depth description of virtual switch configuration is outside the scope of this paper. Your specific configuration depends on physical server and network resources and bandwidth and redundancy requirements. For more information about the networking features of vSphere 4, see the *What's New in VMware vSphere™ 4:Virtual Networking* white paper. At a minimum, you need three distributed virtual switches to support a virtualized SharePoint, one for service console traffic, one for VMotion traffic and one for virtual machine traffic. Note that while this configuration works for smaller production environments and test environments, it offers no redundancy and limited bandwidth for virtual machine traffic.

Hitachi Data Systems testing for this solution used four NICs per ESX host. The first NIC was allocated to the default virtual switch and was used for service console traffic. The second NIC was allocated to a distributed virtual switch and was used for vMotion traffic. The third and fourth NICs were assigned to two separate distributed virtual switches and used for virtual machine traffic. In most cases a production environment uses additional NICs for redundancy or bandwidth requirements and might use VMware features including virtual LANs (VLANs) or NIC teaming.

Conclusion

This white paper describes a solution that enables companies to simplify deployment of Microsoft Office SharePoint Server 2007 in a virtualized environment. It describes a storage solution and computing infrastructure that facilitates centralizing data and scaling the infrastructure as content and users increase.

The Hitachi Adaptable Modular Storage 2000 family provides a reliable, flexible, scalable and cost-effective modular storage system for SharePoint deployments. The 2000 family is ideal for more demanding application requirements and delivers enterprise-class performance, capacity and functionality at a midrange price. By deploying SharePoint on the 2000 family in a VMware environment, IT organizations gain additional advantages, such as increased server utilization resulting in lower costs and increased ease of management.

For more information about the Hitachi Adaptable Modular Storage 2000 family, visit the Hitachi Data Systems Web site or contact your Hitachi sales representative or channel partner.

Hitachi Data Systems Corporation

Corporate Headquarters 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Asia Pacific and Americas 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Europe Headquarters Sefton Park, Stoke Poges, Buckinghamshire SL2 4HD United Kingdom
Contact Information: + 44 (0) 1753 618000 www.hds.com / info.uk@hds.com

AS-036-00 February 2010