

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

Implementation Guide

By Steven Burns and Ken Ewers

February 2010

Summary

Deploying a virtualized Microsoft® SharePoint® Server 2007 environment with VMware vSphere 4 can be a complex endeavor without proper planning and deployment guidance. However, when effectively implemented, the benefits include the reduced storage and server footprints, lower power and cooling costs through consolidation, simplified management of SharePoint Server and SQL Server infrastructure, high availability and easily scalable storage deployments.

This white paper describes the planning and deployment process to help large and enterprise deployments achieve these critical business objectives. It provides planning, installation, and configuration procedures and related information required to successfully deploy a SharePoint Server 2007 farm in a virtualized environment using VMware vSphere 4 on the Hitachi Adaptable Modular Storage 2000 family. It focuses on the storage-related aspects and benefits of implementing a SharePoint 2007 farm in a virtual environment.

The white paper is intended for use by IT personnel responsible for the administration of the SQL Server instances required for a SharePoint environment, server virtualization or 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

Solution Overview	5
Solution Components	7
Hardware.....	7
Software	8
Solution Configuration.....	9
SAN Configuration	9
Storage Configuration	10
VMware Data Center and ESX Server Configuration.....	16
Virtual Machine Creation and Configuration.....	18
Application Configuration	22
Conclusion.....	23

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

Implementation Guide

Deploying a virtualized Microsoft® SharePoint® Server 2007 environment with VMware vSphere 4 can be a complex endeavor without proper planning and deployment guidance. However, when effectively implemented, the benefits of virtualizing a SharePoint Server 2007 farm using VMware vSphere 4 and the Hitachi Adaptable Modular Storage 2000 family outweigh the complexity in many ways, and include the following:

- Reduced storage and server footprints and power and cooling costs through consolidation
- Simplified management of SharePoint Server and SQL Server infrastructure
- High availability and easily scalable storage deployments

This white paper describes the planning and deployment process to help large and enterprise deployments achieve these critical business objectives. It provides planning, installation, and configuration procedures and related information required to successfully deploy a SharePoint Server 2007 farm in a virtualized environment using VMware vSphere 4 on the Hitachi Adaptable Modular Storage 2000 family. It focuses on the storage-related aspects and benefits of implementing a SharePoint 2007 farm in a virtual environment.

The white paper focuses on the storage-related aspects of planning and deploying a SharePoint Server 2007 farm on virtual machines running on VMware ESX 4.0 hosts and is intended for use by IT personnel responsible for the administration of the SQL Server instances required for a SharePoint environment, server virtualization or storage. It assumes familiarity with VMware vSphere 4 and vCenter, Hitachi Storage Navigator Modular 2 software, Microsoft Windows 2008, Microsoft SQL Server 2008, and Microsoft SharePoint Server 2007.

Although the example solution in this paper uses the Hitachi Adaptable Modular Storage 2100, any member of the Hitachi Adaptable Modular Storage 2000 family provides a reliable, flexible and cost-effective storage platform for supporting demanding applications like SQL Server 2008 and SharePoint Server 2007.

Solution Overview

The architecture described in this white paper was developed with ease of management and scalability in mind. The architecture consists of vSphere virtual machines (VMs) running Windows Server 2008, SQL Server 2008 with Service Pack 1, and SharePoint Server 2007 with Service Pack 2 with storage hosted on the Hitachi Adaptable Modular Storage 2100. VMware recommends deploying virtual machines on shared storage like the 2100 to take advantage of VMware features such as VMotion, VMware High Availability (HA), and VMware Distributed Resource Scheduler (DRS).

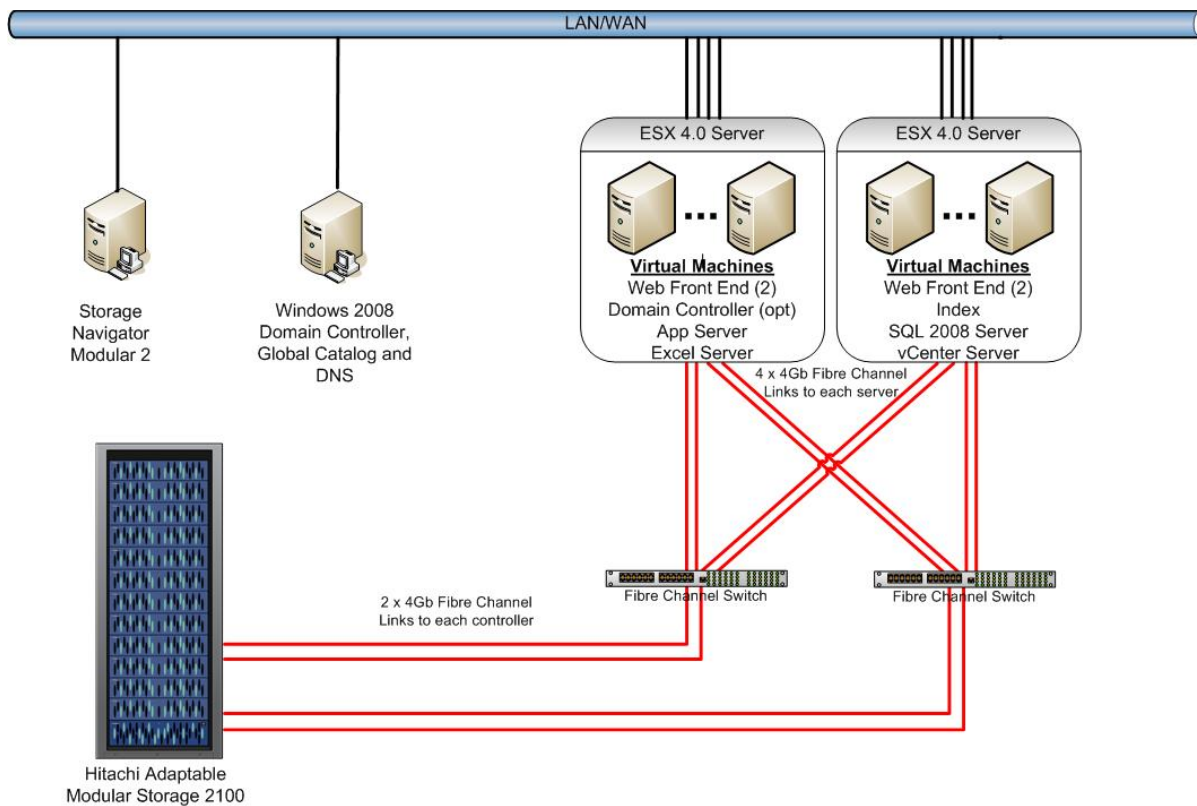
Logical units (LUs) from the 2100 are allocated to the ESX hosts and formatted as VMFS volumes in which virtual disks can be created. The virtual disks (vDisks) are presented to the Windows 2008 operating system and can be partitioned and used in file systems for the operating system boot disks, the SQL Server databases and transaction logs, the SharePoint Index file, and SharePoint Query Catalog files.

The solution described in this white paper supports 1TB of SharePoint content in 10 100GB SharePoint content databases. The storage requirements for the other SQL and SharePoint files are also documented. A SharePoint 2007 implementation consists of three tiers, the Web tier, the application tier and the database tier. The distribution of the virtual machines across the tiers is as follows:

- **Web** — Four Web Front-end servers load-balanced using Microsoft NLB role. The SharePoint query service is enabled on these servers.
- **Application** — Three Application servers. One with the SharePoint Index service enabled, one with the SharePoint Excel service enabled, and one supporting the Shared Service Provider and with the central administration role enabled
- **Database** — One server with SQL 2008 Server installed.

Figure 1 shows the physical layout of the documented solution.

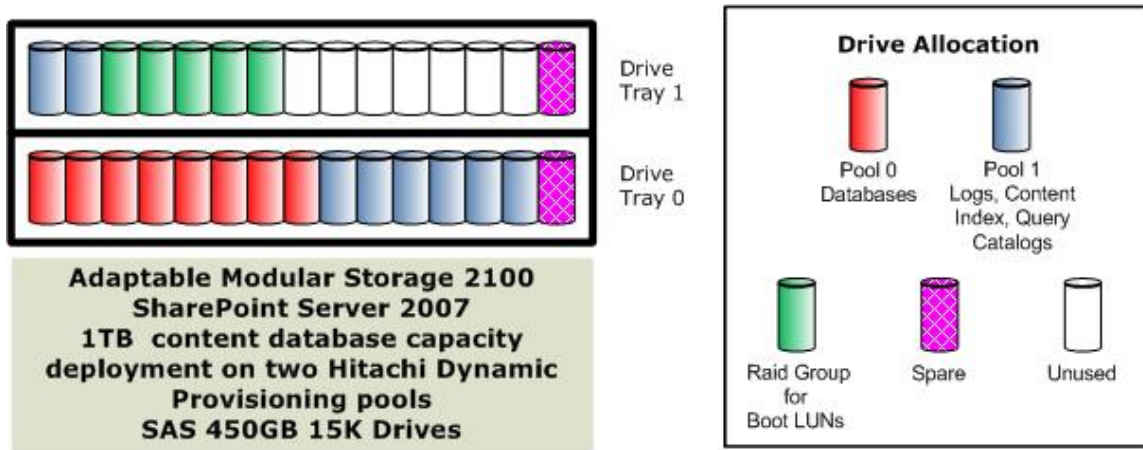
Figure 1. Physical Layout of Solution



This solution uses Hitachi Dynamic Provisioning software for configuring storage. Two Dynamic Provisioning pools were created. One is for the SQL Server database, and tempdb volumes and a separate one is used for the transaction logs volumes, index file, and query catalog files. This gives you the best performance and minimizes hot spots across the RAID groups while maintaining separation of the database and transaction log files. Nevertheless, 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.

Figure 2 shows the storage architecture for the solution. The underlying structure of each Dynamic Provisioning pool is two RAID 1+0 (2D+2D) sets of 450GB 15K drives.

Figure 2. Dynamic Provisioning Environment for SharePoint 2007 Deployment



In addition to the two Dynamic Provisioning pools, one RAID-5 (4D+1P) group stores the OS boot volumes. These are stored on the SAN to simplify migration of virtual machines between ESX hosts if necessary.

For more information about the architecture for this solution and Hitachi Dynamic Provisioning software, see the [Virtualizing Microsoft® Office SharePoint® Server 2007 Using VMware vSphere 4 on the Hitachi Adaptable Modular Storage 2000 Family Reference Architecture](#) white paper.

Solution Components

This section lists the components required for this solution.

Hardware

Table 1 lists the hardware required for this solution. Certain features of the Hitachi Adaptable Modular Storage 2100 used in this solution require specific minimum microcode levels.

Table 1. Hardware Resources for Solution Deployment

Hardware	Configuration
Hitachi Adaptable Modular Storage 2100 storage system	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 Fibre Channel 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 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.

Hitachi Adaptable Modular Storage 2000 Family

This solution was tested on a Hitachi Modular Storage 2100. However, any of the Hitachi Adaptable Modular Storage 2000 family storage systems can be used. Evaluate your capacity and connectivity requirements to determining which member of the 2000 family is best for you.

Hitachi Data Systems testing used a Hitachi Adaptable Modular Storage 2100 with 450GB 15K RPM SAS disks and two controllers with 2GB cache and two ports per controller.

Servers

Use two enterprise-class servers when deploying this SharePoint Server solution. The servers must have enough processor and memory resources to support the SharePoint Server and SQL Server environments and meet vSphere operating requirements. In a production environment, VMware recommends that they both have the same hardware architecture and configuration and processor and memory resources to accommodate a failover situation. Best practice is to not over-commit the ESX host resources (CPU or memory), as this can lead to a bottle neck of resources during peak database utilization times.

For more information about vSphere server hardware requirements, see the [VMware Hardware Compatibility List](#).

Software

Table 2 lists the software used for deploying the tested environment. Note that Windows Server 2008, SQL Server 2008, SharePoint Server 2007 and VMware vCenter 4.0 are installed on one or more virtual machines on the ESX hosts. Note that specific minimum versions of Hitachi Storage Navigator Modular 2 software must be used to manage and monitor certain features of the Hitachi Adaptable Modular Storage 2100.

Table 2. Software Resources for Solution Deployment

<i>Software</i>	<i>Version</i>
VMware ESX	4.0.0, build 164009
VMware vCenter	4.0, build 162856
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 or later**

*Note: The tested deployment used a slipstream build of SharePoint Server 2007 with Service Pack 2.

**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.

VMware vSphere 4 Versions

This solution uses the Enterprise Plus version of vSphere 4. If you use a different version of vSphere 4, lower limits for cores per processor, memory, and vSMP support are available, and you will not be able to use certain features that are not supported by other versions. For more information about vSphere 4 versions, see VMware's [Web site](#) and [Configuration Maximums](#) guide, which lists the VMware supported limits for configuring virtual machines, ESX hosts and vCenter.

Solution Configuration

This section describes the steps required for the implementation of a SharePoint 2007 farm in a vSphere4 environment using a Hitachi Adaptable Modular Storage 2100 for storage.

SAN Configuration

This solution uses two ESX hosts due to processor and memory requirements. Hitachi Data Systems recommends connecting the ESX hosts and the 2100 to two Fibre Channel switches or two separate blades on a Fibre Channel director-class switch for redundancy and high availability.

In addition, Hitachi Data Systems recommends configuring a minimum of two redundant paths from the ESX host to the 2100. Each ESX host requires at least two host bus adapters (HBAs) for high availability. VMware's round-robin load balancing algorithm automatically selects a path by rotating through all available paths, thus balancing the load across all available paths, optimizing IOPS and response time. This solution uses four HBAs per ESX host for additional redundancy and bandwidth.

Configuring the SAN in an ESX environment requires careful consideration of all ESX hosts. For this solution, each ESX host has four HBA ports. The SAN is configured with eight zones, one for each host port. Each zone consists of one host port and one storage port as shown in table x below.

Table 3 lists the connections between the ESX hosts and the storage system ports. Each row in table 3 also defines a separate zone.


Table 3. Path Configuration for ESX Hosts

<i>ESX Host</i>	<i>Host HBA Number</i>	<i>Director Zone Name</i>	<i>Storage System Port</i>	<i>Storage System Host Group</i>
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

Ensure that both the storage and host Fibre Channel ports are configured to point used in your environment.

Figure 4 shows this solution's physical connections between the host HBA ports, the switch ports (or ports on blades of a director) and the controller ports on the 2100.

Figure 4. Fibre Channel Connections



Storage Configuration

The Hitachi Adaptable Modular Storage 2100 allows for easy configuration and deployment of storage solutions through the use of the Hitachi Storage Navigator Modular 2 software and Hitachi Device Manager software. This white paper provides step-by-step instructions for provisioning and managing the storage for your SharePoint environment using the Storage Navigator Modular 2 software GUI. A command-line interface for Storage Navigator Modular 2 software is also available.

Before configuring the 2100, verify that all required licenses are installed.

Configuring the 2100 for a virtualized SharePoint 2007 environment requires the following high-level steps:

1. Create Dynamic Provisioning pools.
2. Create LUs.
3. Configure Fibre Channel ports on the 2100
4. Configure host storage groups on the controller ports
5. (Optional) Assign nicknames to HBA WWNs.
6. Assign LUs to host groups.

Using Hitachi Dynamic Provisioning Software

The Hitachi Adaptable Modular Storage 2100 can be configured to use Hitachi Dynamic Provisioning LUs instead of standard provisioned LUs. In this configuration, all data is distributed across all of the hard disk drives (HDDs) in the Dynamic Provisioning pool. This helps prevent contention for I/O on heavily used LUs by distributing the usage across all HDDs.

Dynamic Provisioning requires the following components:

- Hitachi Adaptable Modular Storage 2000 family microcode version 0870 or higher; microcode version 883/A or higher is required to support pool optimization features
- Storage Navigator Modular 2 software version 7.0 or higher; version 8.3 or higher is required to support pool optimization features
- Hitachi Dynamic Provisioning software license key

Hitachi Dynamic Provisioning software enables you to create a Dynamic Provisioning pool that is made up of one or more RAID groups. A Dynamic Provisioning LU does not consume space from the Dynamic Provisioning pool until the host writes to the Dynamic Provisioning LU.

Creating Dynamic Provisioning Pools

This solution requires two Dynamic Provisioning pools. To review the physical layout of the pools, see Figure 2. One pool is used for all of the SQL databases and the second pool holds the SQL transaction logs, SharePoint index file, and the SharePoint query catalog files. Both pools are created using seven RAID-1+0 (2D+2D) RAID groups. When using Hitachi Dynamic Provisioning software, the creation of the underlying RAID groups is part of the pool creation process, not a separate procedure.

To create Dynamic Provisioning pools using Hitachi Storage Navigator Modular 2 software, follow these steps.

1. Ensure that you have modify privileges on your Storage Navigator Modular 2 software instance.
2. Log in to the Storage Navigator Modular 2 software instance on your data center.
3. Click the **Array Name** link to open the storage system to be used.
4. Expand the **Groups** heading in the storage system's pane and click the **Logical Units** link.

The **Logical Units** pane appears. The right pane has three tabs: **Logical Units**, **RAID Groups** and **DP Pools**.

5. Click the **DP Pools** tab and click the **Create Pool** button.

The Create DP Pool window displays. It has two tabs, **Basic** and **Advanced**.

6. On the **Basic** tab, accept the default value in the **DP Pool** field.

If you prefer a different number scheme, you can change this value.

7. From the **RAID Level** drop-down menu, choose **RAID 1+0**.
8. From the **Combination** drop-down menu, choose **2D+2D**.
9. Accept the default value in the **Number of Drives** field.

This value is determined by your RAID group and drive combination selections. For this solution, the correct value is 4.

10. In the **Drives** section, ensure that the **Automatic Selection** radio button is set.
11. On the **Advanced** tab, review the alerts and thresholds settings.
12. Click **OK**.

A “DP pool created successfully” displays.

13. Click the **Close** button.

The new Dynamic Provisioning pool is now displayed in the list of Dynamic Provisioning pools. The pool consists of a single 2D+2D RAID group at this time. To add capacity, continue with Step 14.

14. Select the check box for the new pool and click **Add Pool Capacity** button.

The Add DP Pool Capacity screen displays.

15. In the **Drives** section, ensure that the **Automatic Selection** radio button is set and click **OK**.

The **Optimize after adding capacity** check box can be left unselected because this pool does not contain any data yet.

A “DP pool capacity added successfully” message displays.

16. Click **Close**.

The pool is now displayed with the new capacity.

17. Repeat Step 14 through Step 17 until the pool has the desired capacity.

18. Repeat these steps to create additional Dynamic Provisioning pools.

For a SharePoint 2007 environment, Hitachi Data Systems recommends using 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.

Creating LUs

In a vSphere 4 environment, the LUs created on the 2100 are used either as VMware datastores to hold the .vmdk files used as virtual disks by the guest virtual machines or as raw device mappings (RDMS) presented directly to a virtual machine. This solution does not use RDMS. For more information about when RDMS might be required see the [Virtualizing Microsoft® Office SharePoint® Server 2007 with VMware vSphere 4 on the Hitachi Adaptable Modular Storage 2000 Family Reference Architecture Guide](#) white paper..

Table 4 lists the LUs that are created for this solution. The capacity of some of the LUs might need to be modified for specific environments. See the RA for more information.

Table 4. LUs Created for Solution

<i>Dynamic Provisioning Pool</i>	<i>LU Number</i>	<i>Purpose</i>	<i>Capacity</i>	<i>Server Assignment</i>
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	0015	WFE2 query catalog	150GB	WFE2
Pool 001	0015	WFE3 query catalog	150GB	WFE3
Pool 001	0015	WFE4 query catalog	150GB	WFE4

• Note: LU 11 stores 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 stores all of the database transaction logs.

To create LUs using Hitachi Storage Navigator Modular 2 software, follow these steps:

1. Expand the **Groups** heading in the storage systems' pane and click the **Logical Units** link.

The **Logical Units** pane appears. The right pane has three tabs: **Logical Units**, **RAID Groups** and **DP Pools**.

2. On the **Logical Units** tab, click the **Create LU** button.

The Create Logical Unit window appears.

3. In the **Type** section, select the **DP Pool** radio button.
4. From the **RAID Group/DP Pool Number** drop-down menu, select the pool number.
5. In the **LUN** field, enter the number that you want assigned to the new LU.

6. In the **Capacity** field, enter the capacity value for the new LU.
7. Click **OK**.

A pop-up window displays indicating that the LU was created successfully.

8. To create additional LUs, click the **Create More LU** button. .

If you are done creating LUs, click **Close**.

Configuring Fibre Channel Ports

In preparation for this step, all ESX HBAs need to be connected to a Fibre Channel switch and appropriately zoned to the storage system ports. Verify that the proper zoning configuration is in place in your environment. If the ports are in use it is important to note that changes to the port's topology setting will cause an interruption of I/O between any host that might be connected to the port.

To configure Fibre Channel ports on the 2100 using Hitachi Storage Navigator Modular 2 software, follow these steps:

1. Expand the **Settings** heading and click the **FC Settings** link.

Check the **Topology** column for the ports that were zoned to connect to the ESX hosts on the SAN. If any of them are not configured to **Point-to-Point**, you must change their configuration.

2. Click the **Edit FC Port** button.
3. Choose **Point-to-Point** from the **Topology** drop-down menu and click **OK**.

A disclaimer displays indicating that the change will cause I/O interruption between any host that is connected to the port at the moment.

4. Click the **Confirm** button and wait a couple seconds for the change to take place.

A message displays stating that you must refresh Storage Navigator Modular 2 software information to see the change. Note that it might also be necessary to reset the ESX host's ports before they can connect to the storage system.

After the connection between the storage system and the ESX host is established, the FC Settings screen shows all modified ports with a **LinkUp(F_Port Connected)** status.

Configuring Host Storage Groups

Enabling host group security enables you to create multiple host groups on a storage port and associate specific WWNs to a host group on a port and a set of LUs within that port. This provides more granularity in cases where more than a single host is connected to the storage port. This also allows servers with different operating systems to attach to the same port by having multiple host groups with different platform settings.

Follow these steps to enable host group security using Hitachi Storage Navigator Modular 2 software:

1. Expand the **Groups** heading and click the **Host Groups** link.
2. Click the **Host Group Security** tab and ensure that port security is enabled on all ports.

If host group security is not enabled follow these steps:

- a. Select all ports to be used in the environment and click the **Change Host Group Security** button.
- b. Select the **Yes** check box and click **OK**.

The **Host Security Group** setting is now enabled on all selected ports.

Assigning Nicknames to HBA WWNs

This procedure is optional, however, assigning a nickname to ESX host HBA's WWNs can facilitate identifying each host and port to ease documenting and troubleshooting your environment. Note that this procedure assumes that the necessary connections and zones exist on the SAN and that the correct settings are configured on the storage ports and host HBAs.

Follow these steps to assign a nickname to an HBA's WWN using Hitachi Storage Navigator Modular 2 software:

1. Expand the **Groups** heading and click the **Host Groups** link.

2. Click the **WWNs** tab.

This tab shows all previously zoned WWNs with each WWN assigned to a Fibre Channel port.

3. Select a port by checking its box in the far left column.

4. Click the **Change Nickname** button.

5. Enter a nickname for the WWN being edited and click **OK**.

A message displays stating the edition of the WWN nickname is complete.

6. Click the **Close** button.

The Storage Navigator Modular 2 software screen automatically refreshes showing the new nickname for the WWN.

Repeat this procedure for the remaining HBA WWNs.

Assigning LUs to a Host Group

Follow these steps to assign LUs to a host group using Hitachi Storage Navigator Modular 2 software:

1. Expand the **Groups** heading and click the **Host Groups** link.

2. Click the **Create Host Group** button.

The **Host Group No.** field can remain the same because Storage Navigator Modular 2 software automatically selects the next available number.

3. In the **Name** field, enter a host group name that identifies the hosts connected to it.

4. Select all of the Fibre Channel ports that are connected to the ESX host.

This ensures that multiple paths to the LUs to be used are available both for high availability and failover purposes.

Note that if the HBA ports are not detected by the storage system, their WWNs are not shown in the **Detected WWNs** window. In that case, it might be necessary to troubleshoot the environment to ensure that the proper settings are in place on both the storage ports and host HBAs.

5. Select the World Wide Names for the host group and click **Add**.

The screen automatically refreshes and shows the added WWNs in the **Selected WWNs** column.

6. Click the **Logical Units** tab.

7. Select the LUs to be added to the host group and click **Add**.

The screen automatically refreshes and shows the added LUs in the **Assigned Logical Units** column.

8. Click the **Options** tab.
9. Select **VMware** from the **Platform** drop-down menu and click **OK**.

Repeat this procedure for each ESX hosts being added to the storage cluster by selecting the appropriate ports, host WWNs and LUs.

VMware Data Center and ESX Server Configuration

A virtual environment has two types of servers, physical (or host), and virtual (or guest). The installation and initial configuration of the ESX 4 server is outside of the scope of this white paper. It is important to note, however, that the physical server must be sized to handle the desired number and type of virtual machines.

Before virtual machines can be created, you must create and configure the VMware datastores and configure VMware networking.

Creating Datastores

Datastores are objects in a VMware data center that are storage locations for virtual machine files. Follow these steps to create a datastore in a VMware data center using vCenter:

1. Use the vSphere client to log in to the vCenter console.
2. Select **Home > Inventory > Datastore**.
3. Select the data center in the left window.
4. Click the **Getting Started** tab.
5. Click the **Add a datastore** link.

The **Add Storage** wizard launches.

6. On the wizard's Select the Host page, click the host where you want to add the datastore and click **Next**.

For this solution, you can choose either host because the datastore will be accessible from both.

7. On the wizard's Select Storage Type page, select the **Disk/LUN** radio button and click **Next**.
8. On the wizard's Select Disk/LUN page, highlight the LU to be used for this datastore and click **Next**.
9. On the wizard's Current Disk Layout page, verify the layout and click **Next**.
10. On the wizard's Properties page, enter the name for the datastore and click **Next**.
11. On the wizard's Disk/LUN – Formatting page, from the **Maximum file size** drop-down menu, select a value for this datastore.

Leave at default of 256GB, Block size:1MB unless you expect larger files (for example, index).

12. On the wizard's Disk/LUN – Formatting page, ensure that the **Maximize capacity** check box is selected and click **Next**.
13. On the wizard's Ready to Complete page, verify that all of the information is correct and click **Finish**.

The task for creating the datastore in the **Recent Tasks** pane appears at the bottom of the vCenter window. This takes a few seconds to complete.

Configuring Multipathing

By default a datastore is configured with the **Fixed** multipathing algorithm. This solution uses the round robin multipathing algorithm. Follow these steps to configure multipathing for a datastore using vCenter.

1. Use the vSphere client to log in to the vCenter console.
2. Select **Home > Inventory > Host and Clusters** in vCenter.
3. Highlight a host in the left window.

A set of tabs appears in the right pane.

4. Click the **Configuration** tab.
5. In the **Hardware** pane on the left side, highlight **Storage**.
6. Click the **Devices** button.
7. In the **Devices** box, highlight the device (LU) for which you want to configure multipathing.
8. Click the **Manage Paths** link.

A pop-up window displays.

9. From the **Path Selection** drop-down menu, choose **Round Robin (VMware)** and click **Close**.

Two tasks run in the Recent Tasks window at the bottom of the vCenter screen.

10. When both tasks complete, click the **Manage Paths** link again and verify that all of the paths show **Active (I/O)** instead of just **Active**.

VMware Network Configuration

This solution uses four NICs per ESX host. All NICs are gigabit Ethernet. These are allocated to the following kinds of traffic:

- **Service console** — One NIC is used to handle traffic for the service console on the ESX host.
- **Virtual machine** — Two NICs are used to handle all network traffic for the VMs.
- **vMotion** — One NIC is used to support moving VMs between hosts, storage or both.

Each of the NICs is assigned to either a virtual switch or a distributed virtual switch. A virtual switch is configured at the ESX host level and a distributed switch is configured at the VMware data center level. This solution uses a virtual switch for the service console and virtual distributed switches for VM and vMotion traffic.

To add a distributed virtual switch to a VMware data center follow these steps:

1. Use the vSphere client to log in to the vCenter console.
2. Select **Home > Inventory > Networking** in vCenter.
3. Right-click on the data center in the left window and choose **New vNetwork Distributed Switch** from the drop-down menu.

The **Create vNetwork Distributed Switch** wizard launches.

4. On the wizard's General Properties page, enter a name for the distributed switch in the Name field, choose a value in the **Number of dvUplink ports** field and click **Next**.
5. On the wizard's Ready to Complete page, deselect the **Automatically create a default port group** check box, verify the settings and click **Finish**.

6. After the switch is created, right-click the switch in the left window and click **New Port Group** from the drop-down menu.

The **Create Distributed Virtual Port Group** wizard launches.

7. On the wizard's Properties page, enter a descriptive name for the port group in the **Name** field, leave all other settings at default values and click **Next**.
8. On the wizard's Ready to Complete page, verify the settings and click **Finish**.

This completes the network configuration for a switch that will be used for VM traffic. To configure a switch that will be used for vMotion traffic, continue with Step 9.

9. Select **Home > Inventory > Hosts and Clusters**.
10. Highlight a host in the left window and click the **Configuration** tab in the right window.
11. In the **Hardware** pane, highlight **Networking**.
12. Click the **Distributed Virtual Switch** button.
13. Click the **Manage Virtual Adapters** link for the switch you want to configure.

A Manager Virtual Adapters window displays.

14. Click **Add** in the menu bar.

The **Add Virtual Adapters** wizard launches.

15. On the wizard's Creation Type page, select the **New virtual adapter** radio button and click **Next**.
16. On the wizard's Virtual Adapter Type page, select the **VMkernel radio button** and click **Next**.
17. On the wizard's Connection Settings page, select the **select port group** radio button and select the correct port group from the drop-down menu that becomes active, select the **Use this virtual adapter for vMotion** check box and click **Next**.
18. On the wizard's VMkernel – IP Connection Settings page, enter an IP address and subnet mask for your environment and click **Next**.
19. On the wizard's Ready to Complete page, verify the settings and click **Finish**.

The process to add a virtual adapter for vMotion must be repeated for each host.

Virtual Machine Creation and Configuration

Templates shorten the time required to create new virtual machines. This solution uses 64-bit virtual machines.

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

Use the thick format the .vmdk file for all hard disks that are added to your virtual machines. This is the default and is used unless the thin provisioning or fault tolerance options are selected the hard disks are created.

For more information about virtual machine templates, see VMware's support [Web site](#).

Creating a Virtual Machine Template

To create a virtual machine template using vCenter, follow these steps:

1. Use the vSphere client to log in to the vCenter console.
2. Select **Home > Inventory > VMs and Templates**.
3. Right-click a data center in the left pane and select **New virtual machine** from the drop-down menu.

The **Create New Virtual Machine** wizard launches.

4. On the wizard's Configuration page, select the **Custom** radio button.

Custom allows you to select the number of vCPUs and the amount of memory allocated to the VM.

5. On the wizard's Name and Location page, assign a name to the VM, click an inventory location if more than one data center is managed by the vCenter instance and assign a descriptive name.
6. On the wizard's Host/Cluster page, select the ESX host where you want this VM to run.

Because you're creating a template, it doesn't matter which host you choose.

7. On the wizard's Datastore page, select the datastore where you want the OS boot disk to be stored.

This solution uses the datastore that uses the LU created on the wizard's Virtual Machine Version page, select the **Virtual Machine version: 7** radio button.

8. On the wizard's Guest Operating Systems page, select the **Microsoft Windows** radio button and choose **Microsoft Windows Server 2008 (64-bit)** from the **Version** drop-down menu.
9. On the wizard's CPU page, set the number of virtual processors to 2.
10. On the wizard's Memory page, set the memory to 2GB.
11. On the wizard's Network page, set the number of NICs to connect to 1, set the network for NIC 1 to the network configured for VM traffic, set the adapter to **E1000**, select the **Connect at Power On** check box and click **Next**.

12. On the wizard's SCSI Controller page, select the **LSI Logic SAS** radio button.

13. On the wizard's Select a Disk page, select the **Create a new virtual disk** radio button.

14. On the wizard's Create a Disk page, set the capacity to 40GB, **leave the two check boxes in the Disk Provisioning** section unchecked and select the **Store with the virtual machine** radio button in the **Location** section.

15. On the wizard's Ready to Complete page, verify the configuration and click **Finish**.

16. After the virtual machine is created, install Windows Server 2008 (64-bit).

17. Add all patches and any additional software needed for the VM guest operating system.

18. Install VMtools on the VM.

19. Shut down the VM, right-click the VM in vCenter and select **Convert to Template** from the **Template** menu.

Deploying Virtual Machines from a Template

Virtual machines can be deployed using the template created in the previous procedure. Be sure to use the customization wizard to generate a new SID. To deploy a new virtual machine from a template, follow these steps:

1. Use the vSphere client to log in to the vCenter console.
2. Select **Home > Inventory > VMs and Templates**.
3. Right-click the template in the left pane, and select **Deploy to a new virtual machine** from the drop-down menu.

The **Deploy Template** wizard launches.

4. On the wizard's Name and Location page, assign a name to the VM and select an inventory location if more than one data center is managed by the vCenter instance.
5. On the wizard's Host/Cluster page, click the ESX host where you want this VM to run.
6. On the wizard's Datastore page, select the datastore where you want the OS boot disk to be stored.

This solution uses the datastore that uses the LU created on the RAID-5 (4D+1P) group

7. On the wizard's Disk Format page select the **Same format as source** radio button.
8. On the wizard's Disk Customization page, select the **Customize using the Customization Wizard** radio button or if you have already created and saved a customization specification select the **Customize using an existing customization specification** radio button
9. Click **Next** and complete the customization screens.

Make sure that you select the **Generate New Security ID (SID)** check box on the Operating System Options page.

If you are creating a new customization specification, you can name and save it from the Save Specification page.

10. On the wizard's Ready to Complete page for the customization, verify that the configuration is correct and click **Finish**.
11. On the wizard's Ready to Complete page for deploying the VM, verify that the configuration is correct and click **Finish**.
12. After the VM is created, right-click the VM in the left pane of vCenter and click **Edit Settings** from the drop-down menu.

The Virtual Machine Properties window displays. You can edit the settings for number of processors and amount of memory depending on the requirements of the specific VM. You can also add additional SCSI controllers and hard disks as required.

Settings for SharePoint VMs

When creating the SharePoint VMs from the template, the following settings for number of vCPUs, amount of memory and hard drive configuration were used for this solution.

- **WFE VMs** — This solution uses four WFE VMs, divided evenly between the two ESX hosts for high availability. These were created from the VM template but the following settings were changed:
 - Increase memory from 2GB to 4GB.
 - Configure a second SCSI controller.
 - Attach an additional drive to the second controller. This drive is where the query catalog is stored. The vDisk for this drive is stored on the datastore created for this purpose.
- **Index server** — The index server is an application server with the search service enabled for indexing purposes. It is listed separately here due to the difference in resource requirements. Create this server from the template but make the following changes to the settings.
 - Increase number of vCPUs from two to four.
 - Increase the memory from 2GB to 8GB.
 - Configure a second SCSI controller.
 - Attach an additional drive to the second controller. This drive is where the content index is stored. The vDisk for this drive is stored on the datastore created for this purpose. 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 indexed. 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 will be smaller.
- **Application servers** — This includes servers running the Excel service. Create these from the VM template. The only modification that needs to be made to the settings for these servers is to increase the memory from 2GB to 4GB.
- **SQL server** – Create this server from the VM template. Make the following modifications to the settings:
 - Increase number of vCPUs from two to four.
 - Increase the memory from 2GB to 16GB.
 - Configure two additional SCSI controllers.
 - Add the SQL database and log drives to the second and third controllers.

While VMware recommends allocating two vCPUs and 6GB of memory and scaling up as necessary, additional memory increases caching efficiency of the SQL server, which decreases the number of IOPs for the storage and increase performance. For that reason, Hitachi Data Systems recommends allocating 12GB to 16GB of memory.

Your requirements might vary depending upon the specific environment. For more information about adapting this solution to your environment, see the [Virtualizing Microsoft® Office SharePoint® Server 2007 with VMware vSphere 4 on the Hitachi Adaptable Modular Storage 2000 Family Reference Architecture Guide](#) white paper.

Application Configuration

Installation procedures for SQL Server 2008 and SharePoint 2007 are outside the scope of this document. This section describes tasks or settings that differ from standard installation or configuration.

Installing SQL Server 2008

Install SQL Server 2008 using standard procedures. On the Windows server where SQL is installed, create a single volume on each of the disks and mount these using mount points to the server. Name these volumes and mount points to reflect their purpose, for example Content DB1, logs, tempdb, and so on.

Configuring SQL Server 2008

When configuring SQL Server 2008 for this solution, do the following:

- Move the tempdb files to the tempdb volume on the server.
- Create multiple data files for tempdb, as follows:
 - Create one data file for each CPU core on the server.
 - Pre-size these files.
- Pre-build and pre-allocate space for each of the following:
 - Each of the SharePoint content databases
 - SharePoint search database
- Pre-size log files in 4GB or 8GB segments to avoid fragmentation.

Installing SharePoint Server 2007

Install SharePoint Server 2007 using standard procedures. Make sure that all files created on the SQL server during installation, such as the configuration database and logs, are placed in the correct location on the SQL server.

For the Web front end servers and the index server, create a single volume on the second disk to store the query catalog (on the WFEs) or the content index (on the index server).

Initially Configuring SharePoint Server 2007

Verify that all database files and logs are placed in the correct locations on the SQL server.

When configuring the servers with the index or query services enabled store the content index or the query catalog on the volume allocated for this purpose.

Datastore and .vmdk Considerations

When using a datastore for a single .vmdk, which is done with many of the datastores in this solution, alarms are triggered in vCenter due to the perceived lack of free space on the datastore. For example, if the LU presented from the 2100 is 100GB in size and you create a .vmdk file that is 100GB in size on the datastore that uses that LU, warning alarms are triggered in VMware vCenter if the alarm definitions are left at default values. You can avoid these alarms in one of two ways:

- Disabled the alarm in vCenter. The specific alarm is named Datastore usage on disk. Because the alarms are defined at the vCenter level, disabling this alarm disables it for all datastores in the environment.
- Create larger LUs initially by over-provisioning the storage in an Dynamic Provisioning pool. The vDisk file and the NTFS volume are still the same initial size (100GB in this example), but the datastore can be large enough to avoid the alarms. Because the storage capacity is being over-provisioned, it is very important that you monitor pool usage on the 2100, because a risk exists that additional space could be allocated on the datastore without physical storage to support it. It is important that the person responsible for managing the VMware datastores and the virtual machine vDisk files understands that the need to verify that the physical storage capacity exists before enlarging the size of a vDisk file or allocating additional vDisks.

Conclusion

This white paper provides deployment instructions for a solution that enables companies to simplify deployment of Microsoft Office SharePoint Server 2007 in a virtualized environment. 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.



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

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

All other trademarks, service marks and company names mentioned in this document or Web site are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems. This document describes some capabilities that are conditioned on a maintenance contract with Hitachi Data Systems being in effect and that may be configuration dependent, and features that may not be currently available. Contact your local Hitachi Data Systems sales office for information on feature and product availability.

© Hitachi Data Systems Corporation 2010. All Rights Reserved.

AS-037-00 February 2010