

Implementing SharePoint 2007 with the Hitachi Adaptable Modular Storage 2000 Family

Planning and Deployment Guide

By Art LaMountain

October 2009



Summary

Paperless collaboration and content management systems are increasingly popular and indispensable tools that large and enterprise organizations rely on to deal with explosive growth in digital content, increasing needs to facilitate collaboration among employees, and stringent demands by regulators for thorough and quick access to data. Microsoft® Office SharePoint Server 2007 provides these much-needed collaboration, business process and content management capabilities.

The solution for SharePoint deployments described in this white paper uses a building block architecture and the Hitachi Adaptable Modular Storage 2000 family. It allows IT administrators to optimize the 2000 family to provide a common interface to complex SharePoint deployments. The building block architecture designed around the 2000 family is scalable, highly available and easy to manage.

This white paper provides calculations on how to size a SharePoint environment, step-by-step instructions for configuring and deploying the 2000 family storage system, and guidance on validating a new SharePoint deployment.



Contributors

The information included in this document represents the expertise, feedback and suggestions of a number of skilled practitioners. The author recognizes and sincerely thanks the following contributors and reviewers of this document:

- Rick Andersen
- Steven Burns
- Alan Davey
- Eduardo Freitas
- Ron-An Lee
- Larry Meese
- Lisa Pampuch



Table of Contents

Solution Overview	1
Planning	3
Hitachi Adaptable Modular Storage 2000 Family	3
Storage Building Blocks	3
Storage Sizing and Design.....	4
Deployment.....	7
Configuring Storage	7
Server Preparation.....	10
Server Requirements	11
SQL 2008 Installation.....	12
SharePoint Installation	13
Validation	14
SQLIOSim	14
SQLIO	15
Iometer.....	15
SharePoint 2007 Test Data Population Tool	15
Visual Studio Team System Test Edition	15
Confirming SharePoint 2007 Functionality	15
Hitachi Storage Management Software.....	16



Implementing SharePoint Server 2007 with the Hitachi Adaptable Modular Storage 2000 Family

Planning and Deployment Guide

Paperless collaboration and content management systems are increasingly popular and indispensable tools that large and enterprise organizations rely on to deal with explosive growth in digital content, increasing needs to facilitate collaboration among employees, and stringent demands by regulators for thorough and quick access to data. Microsoft Office SharePoint Server 2007 provides these much-needed collaboration, business process and content management capabilities. Microsoft Windows SharePoint Services 3.0, a component of Microsoft Windows Server, provides the framework for easy access to information and people. However, creating a robust SharePoint deployment that is scalable, highly available, and easy to manage can be a complex and costly endeavor. The complexities of a large-scale SharePoint 2007 implementation are simplified with proper storage area network (SAN) storage management and functionality.

The solution for SharePoint deployments described in this white paper uses a building block architecture and the Hitachi Adaptable Modular Storage 2000 family. It allows IT administrators to optimize the 2000 family to provide a common interface to complex SharePoint deployments. The building block architecture designed around the 2000 family is scalable, highly available and easy to manage.

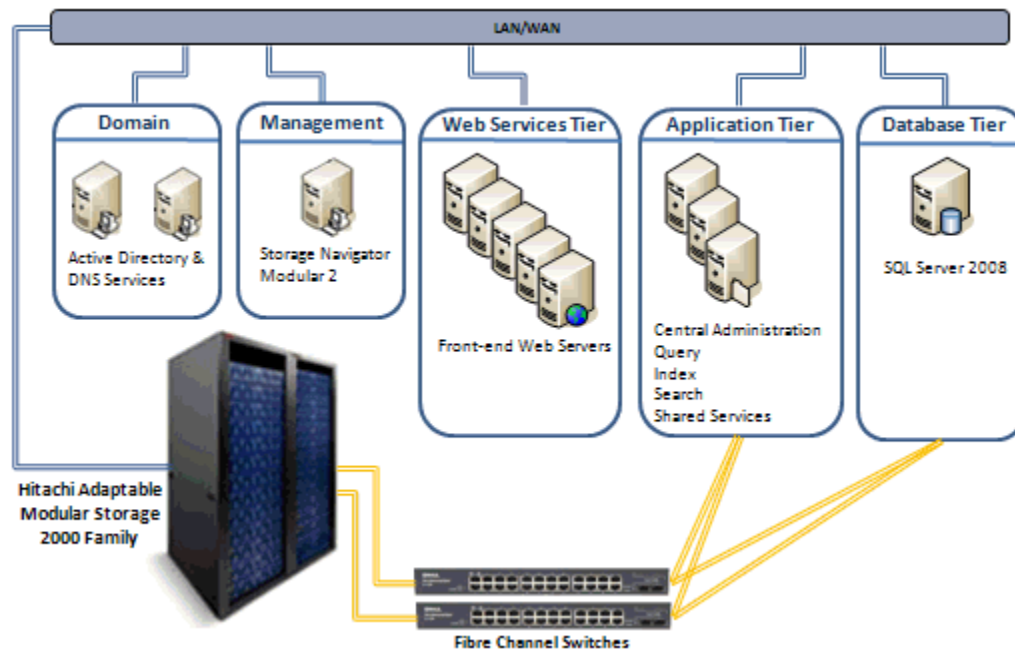
This white paper provides calculations on how to size a SharePoint environment, step-by-step instructions for configuring and deploying the 2000 family storage system, and guidance on validating a new SharePoint deployment. It is written for IT administrators in large and enterprise organizations who need reference information on how to deploy a SharePoint 2007 enterprise environment using the Hitachi Adaptable Modular Storage 2000 family. It assumes a basic level of storage, SAN, SharePoint and Microsoft SQL Server 2008 knowledge.

Solution Overview

Whether you're implementing a single server or a large farm, proper storage planning and design are critical for a successful SharePoint deployment. Hitachi Data Systems testing used a single SQL 2008 server, separated Application servers for Query, Search SSP and Central SharePoint Administration, and multiple front-end Web servers using Network Load Balancing (NLB). The overall environment includes Active Directory and DNS services and a separate server to manage storage using Hitachi Storage Navigator Modular 2 software. All application and SQL servers connected to the storage system include dual HBAs to allow for alternate paths.

Figure 1 illustrates the tiers and windows infrastructure for all deployments.

Figure 1. Solution Topology



A standard deployment, whether in a development, staging or production environment, needs to contain the basic tiers that provide the SharePoint framework, as described in Table 1. For more information, see the Microsoft TechNet article "[Design content deployment topology.](#)" The Hitachi Data Systems lab environment is based specifically on the production environment separating roles and tiers by individual servers.

Table 1. SharePoint Environments

<i>Environment</i>	<i>Description</i>
Development	Houses creation of site collections and content organization to determine best practices and testing to ensure everything flows and works properly. This implementation is often less robust in comparison to staging or production due mainly to the interest of functionality rather than performance.
Staging	Houses a copy of the production site collections, allowing review and testing prior to publishing from development testing. This environment closely resembles the complete roles and tiers of production but still not necessarily the complete storage architecture because failover and performance are not always as important as the true live production environment.
Production	Houses site collection that presents content, usually with tightened security.

Planning

Planning a SharePoint deployment involves understanding end to end (storage to user) requirements to prepare the collaborative environment. This includes storage building blocks, and design considerations including the unique requirements of each type of database, plus storage map points for items such as index, to the SharePoint Farm.

Hitachi Adaptable Modular Storage 2000 Family

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.

For more information, see the [Hitachi Adaptable Modular Storage 2000 Family overview brochure](#).

Logical Units

On a midrange system, when space is carved out of a RAID group and made into a volume, it is then known as a logical unit (LU). The LU is mapped to a host port for use by a server by associating the World Wide Name (WWN) of the HBA port to the storage Fibre Channel port.

Number of LUs per RAID Group

When configuring a midrange storage system, one or more LUs can be created per RAID group, but the goal should be to clearly understand what percentage of that group's overall capacity will contain active data. In the case where multiple hosts attempt to simultaneously use LUs that share the same physical disks in an attempt to fully utilize capacity, seek and rotational latency might be performance-limiting factors. In attempting to maximize utilization, RAID groups should contain both active and less frequently used LUs. This is true of all physical disks regardless of size, RAID level and physical characteristics. It is also true that, if many small LUs are carved out of a single RAID group, their simultaneous use creates maximum seek times on each disk, reducing the maximum sustainable small block random IOPS rate to the disk's minimum.

LU Management and Controller I/O Management

On nearly every midrange storage system from any vendor, the individual LUs are tightly bound to an owning controller. This is because no global sharing exists between the controllers of either the data or its metadata. Each controller is independently responsible for managing these two objects. All LUs are accessible regardless of the physical port or the server from which the access is requested.

The 2000 family's active-active symmetric front-end design allows any host port to access any LU managed by either controller and allows the use of host-based path management and load balancing such as Hitachi Dynamic Link Manager software or MPIO with minimal penalty.

Storage Building Blocks

A storage building block is made up of database locations, configuration of RAID technology and mapping of LUs. These are important areas of review for building the storage environment for SharePoint. Other important considerations include the division of the logical configurations and placement and separation of each type database files, including tempdb and transaction logs. For more information, see the [Implementing SharePoint Server 2007 with the Hitachi Adaptable Modular Storage 2000 Family Best Practices and Reference Architecture Guide](#) white paper.

Database server configuration affects the performance of SharePoint Server 2007. Separate each type of data onto individual spindle sets, with the appropriate RAID level, spindles and optimization for each data type, as shown in Table 2 and Figure 2. These LUs are created on the 2000 family storage system and mapped to the SQL Server 2008 computer to place appropriate targets for installation.


Table 2. RAID Levels and Optimization Characteristics by Database Type and SAN-attached Storage

Type	RAID Level	Optimization
tempdb	RAID-1+0 (2+2)	Write
Transaction logs	RAID-1 (1+1) *	Write
Index	RAID-1+0 (2+2)	Read/write
Search property store (search database)	RAID-1+0 (2+2)	Read/write
Configuration database	RAID-1+0 (2+2)	Read
Content databases	RAID-1+0 (2+2)	Read

* For cost-sensitive environments, use RAID-5 for transaction logs rather than RAID-1+0. In some cases, RAID-5 performs as well as RAID-1+0 on the 2000 family for sequential workloads like Microsoft SQL Server log files. Consult your SQL database administrator about all storage design related configurations.

Figure 2 illustrates storage definitions and separation of LUs as defined by database type and disk RAID type.

Figure 2. Initial Configuration of RAID Groups and LUs by Database Type



Storage Sizing and Design

When designing your SharePoint environment, carefully consider database design and overall performance, and after your environment is implemented, regularly monitor it to ensure the deployment remains healthy and continues to meet your needs.

Database File Considerations

When deploying SQL Server in support of SharePoint databases, it is important to have an overall understanding of the key file types that are essential for a database. Understanding each file type's workload and database type and size enables both storage and database administrators (DBAs) to establish storage requirements for the SQL Server and SharePoint environment.

Configuration Database

The configuration database is used to manage the farm settings. Each farm has one configuration database and it generally never grows past 1.5GB in size.

Content Database

The content databases are used to store content. Your SharePoint farm can have many different content databases. If versioning is used for documents, a copy of each version is stored in the database. Many factors determine and limit the amount and size of the content databases. If versioning is used for documents, a copy of each version is stored in the database. To minimize limitations and enhance performance, it's important to manage hierarchy within your farm.

Consider these factors when estimating the initial size of your content database:

- To estimate the size of stored content in your SQL database, multiply your initial content estimate by 1.3.
- If you plan to use versioning for documents, remember that a copy of each version is stored in the database, so increase your initial estimate accordingly.

For example, if you have a single Web application with two site collections and 20 nested Web sites within those collections, each site can store an estimated 1,000 files. If each stored file is estimated to be an average of 5MB, the resulting storage estimate follows this calculation:

$$20 \times 1,000 \times 5\text{MB} = 100,000\text{MB (100GB)}$$

Microsoft recommends that content databases per site collection not exceed 100GB. When your content databases approach this limit, divide them into multiple databases. The limitation is at the site collection level, which can only handle a single database. When nearing threshold limits, creating new site collections with unique content databases helps keep database sizes within recommended specifications.

When planning a SharePoint deployment, it's crucial to account for future growth. Microsoft recommends doubling your initial data estimate.

Table 3 describes guidelines using the hierarchy recommended by Microsoft.

Table 3. Guidelines by Object

<i>Object</i>	<i>Guidelines</i>
Content databases	Maximum of 100 databases per Web application.
Site collections	50,000 collections per content database. 150,000 collections per Web application.
Web sites	Maximum of 125 sites with 2,000 sub sites (no greater than 250,000 combined sites per site collection). Limit sub sites to 2,000 per Web site.
Documents	5 million per library. Document file size is restricted by system to 2GB and by default to 50MB.

The guidelines described in Table 3 are largely theoretical and depend on factors like these:

- Database server performance
- Web server performance including network load balancing and quantity of machines serving users accesses
- Network bandwidth between the Web servers and the database server

For more information, see the Microsoft TechNet article "[Plan for software boundaries \(Office SharePoint Server\).](#)"

Index

Index and search roles can be affected by many factors, including size and quantity of data, number of users, amount of server roles and more. Each SharePoint farm requires its own unique review after each sizing determination is considered. Use the following formula to calculate how much storage is required for content index:

$$\text{GB of disk space required} = \text{Total_Corpus_Size (in GB)} \times \text{File_Size_Modifier} \times 2.85$$

Note:

- Total_Corpus_Size is the amount of physical file content stored in content databases.
- File_Size_Modifier is based on the average size of the files in your content databases:
 - **1.0** — Contains small files (average file size is 1KB).
 - **0.12** — Contains moderate files (average file size is 10KB).
 - **0.05** — Contains large files (average file size is 100KB or larger).

For example, for a content size of approximately 700GB that primarily contains files averaging 100KB in size, use the following calculation:

$$700\text{GB} \times 0.05 = 35\text{GB estimated size}$$

Next, multiply the estimated size of the index file by 2.85:

$$35\text{GB} \times 2.85 = 100\text{GB size requirement}$$

Search Database

The search database stores metadata and search history information. This database normally requires more space than the index disk space. Use the following formula to calculate how much disk space you need for the search database:

$$\text{GB of disk space required} = \text{Total_Corpus_Size (in GB)} \times \text{File_Size_Modifier} \times 4$$

Note:

File_Size_Modifier is based on the average size of the files in your corpus for complete environment or Total_Corpus_Size:

- **1.0** — Contains small files (average file size is 1KB).
- **0.12** — Contains moderate files (average file size is 10KB).
- **0.05** — Contains large files (average file size is 100KB or larger).

For example, for a content size of approximately 150GB that primarily contains small files averaging 1KB in size, use the following calculation:

$$150\text{GB} \times 1.0 = 150\text{GB}$$

Then multiply the estimated size of the index file by 4:

$$150\text{GB} \times 4 = 600\text{GB}$$

Key Considerations

- Place database and log files on physically separate RAID groups.
- Place log files on RAID-1 or RAID-1+0, depending on their capacity and performance requirements.
- Select the appropriate RAID type for database files based on performance, cost and availability requirements. RAID-1+0 provides the best performance and availability, but at a higher cost. RAID-5 can provide good performance with slightly lower availability and at a lower cost. RAID-6 can also be considered if the higher availability is important enough to justify reduced write performance due to the additional RAID-6 write overhead.
- For cost-sensitive environments, use RAID-5 for transaction logs rather than RAID-1+0. In some cases, RAID-5 performs as well as RAID-1+0 on the 2000 family for sequential workloads like Microsoft SQL Server log files.
- Create one tempdb file per CPU core (per SQL instance) and make all files equal in size. For more information about using multiple files with a database configuration, see the Microsoft TechNet article "[Optimizing tempdb Performance.](#)"
- Place tempdb files on a separate RAID-1+0 group from the database and log files with the fastest hard disk drives available. However, with caution, tempdb files can reside on the same RAID group as the database files. To simplify monitoring, always place tempdb files in a dedicated LU, regardless of RAID group placement. tempdb LU usage and growth must also be evaluated prior to deployment and monitored regularly.

Deployment

Deploying a SharePoint environment includes many steps, from configuring servers to developing the actual site architecture. Detailed instructions about setting up the environment are out of the scope of this document, which focuses on the storage-related aspects of SharePoint deployment. For more information, see the deployment articles available on Microsoft's [TechNet Web site](#).

Configuring Storage

The Hitachi Adaptable Modular Storage 2000 family allows for easy configuration and deployment of storage solutions using Hitachi Storage Navigator Modular 2 and Hitachi Device Manager software. This guide provides step-by-step instructions on how to provision and manage the storage for your SQL Server environment through using Storage Navigator Modular 2 software's graphical user interface. A command-line interface is also available for Storage Navigator Modular 2 software that can be used to accomplish the same tasks.

Configure the following settings using Storage Navigator Modular 2 software:

- Enable Host Group Security for all ports
- Enable LU Mapping mode

For more information about how to set these options, see the Hitachi Storage Navigator Modular 2 software online help.

Creating RAID Groups

Follow these steps to create RAID groups using Hitachi Storage Navigator Modular 2 software:

1. Log in to the Storage Navigator Modular 2 software instance on your data center with modify privileges.
2. Click the **Array Name** link to open the storage system to be used for the SQL Server environment.
3. Expand the **Groups** heading in the array's pane window and click the **Logical Units** link.

Three tabs appear in the right window pane: **Logical Units**, **RAID Groups** and **DP Pools**.

4. Select the **RAID Groups** tab and click the **Create RG** button.

5. Select **RAID1+0** RAID level and the **2D+2D** combination.

This is the RAID group for the database and tempdb LUs.

6. Ensure that the **Number of parity groups** is set to 1.

This is Hitachi Data System best practice recommendation. Increasing the number of parity groups above the default of 1 allows the creation of large LUs across multiple RAID groups. Hitachi Data Systems does not recommend creating LUs that span across multiple RAID groups. Instead, use the Hitachi Dynamic Provisioning software to create LUs that automatically span across multiple hard drives. For more information, see the [Hitachi Dynamic Processing software datasheet](#).

Note that Hitachi Dynamic Provisioning software can be used to provide virtual storage capacity to reduce administration costs by reducing the time to provision new storage, but also improves application availability by reducing the downtime needed for storage provisioning. However, the building block architecture described in this solution uses traditional RAID group design.

7. Select the **Automatic Selection** radio button.

If you have different types of drives installed in the storage system, for example, SATA and serial-attached SCSI (SAS), or different capacities, be sure you select the correct type so that the Storage Navigator Modular 2 software can select the appropriate hard drive type.

Enabling automatic selection means that Storage Navigator Modular 2 software selects the next available drives of the type shown in the box. Hitachi Data Systems recommends enabling this setting.

8. Click **OK**.

A message appears when the RAID group is successfully created.

Repeat these steps to create the all RAID groups that will host database and transaction log LUs.

Creating LUs

Follow these steps to create LUs using Hitachi Storage Navigator Modular 2 software:

1. Expand the **Groups** heading in the array's pane window and click the **Logical Units** link.
2. Select the **Logical Units** tab and click the RAID group where you want to create a LU.

The Summary window opens in the right pane.

3. Click the **Create LU** button.
4. Change the **Capacity** information from RG ALL to GB, and enter the desired LU size.

In this example, the first LU to be created on the RAID group contains space for the first set of content databases. Remember that while these are often targeted as 100GB databases, 20 percent more capacity is being allocated for future or unplanned growth as per Microsoft's best practice recommendations.

5. Click **OK**.

After LU creation is complete, a success message appears. The newly created LUs begin a quick formatting process that makes them immediately available for use.

Repeat these steps to create additional LUs, followed by the creation of the transaction log LU on its own RG.

Changing Fibre Channel Port Settings

In preparation for this step, all Application and SQL servers HBAs must be connected to a Fibre Channel switch and appropriately zoned to the storage system ports.

To ensure the proper zoning configuration is established on your environment, contact your SAN administrator. It is important to ensure the correct port is being edited, because changes to the port's topology setting cause I/O interruption between any host that might be connected to the port.

To ensure that your storage Fibre Channel ports are configured for the correct topology follow these steps:

1. Expand the **Settings** heading and click the **FC Settings** link.
2. Click the **Edit FC Port** button.
3. Change the Topology setting to **Point-to-Point** and click the **OK** button.

A disclaimer appears indicating that the change will cause I/O interruption for any host that is connected to the port.

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

After the change is complete, a message displays indicating that you must refresh Hitachi Storage Navigator Modular 2 software's information to the changes.

After the connection between the storage system and the server is established, the Fibre Channel Settings screens shows all modified ports with a **LinkUp (F_Port Connected)** status.

Enabling Host Group Security

Enabling Host Group Security allows you to associate a specific WWN to a specific ports and to a set of LUs within that port. This enables you to have more granularity in cases where more than a single host is connected to the storage port.

Follow these steps to enable Host Group Security using Hitachi Storage Navigator Modular 2 software:

1. Expand the **Groups** heading and click on the **Host Groups** link.
2. Click the **Host Group Security** tab and ensure all ports have the port security enabled.
3. If host security is not enabled, select all ports to be used in the environment and click the **Change Host Group Security** button.
4. Select the **Yes** checkbox and click **OK**.

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

Assigning Nicknames to HBA's WWNs

Assigning nicknames to your host HBA's WWNs can facilitate identification of each host and port for documentation and troubleshooting.

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 chooses the next available number.

3. Change the **Name** of the host group to something that identifies the hosts connected to it.

4. Select all of the Fibre Channel ports that are connected to the particular application or SQL Server host.

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

Note that if the HBA ports are not detected by the storage system, their WWNs cannot be shown in the **Detected WWNs** window. In that case, it might be necessary to troubleshoot the environment to ensure that the proper HBA connections are discovered.

5. Select the WWNs for the host group and click **Add**.

The screen automatically refreshes and shows the added WWNs on 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 on the **Assigned Logical Units** column.

8. Click the **OK** button.

Server Preparation

The first step in deploying SharePoint is installation of all prerequisite Windows components noted in this section and then configuring the environment.

Every server in your farm requires these components:

- Windows Server 2008 (64 bit)
- Internet Information Services (IIS) 7.0
- Microsoft .NET ASP 2.0
- NET Framework version 3.5
- Microsoft installer 4.5

Keep the following performance and configuration requirements in mind as you plan your SharePoint deployment:

- SharePoint Server 2007 requires the following components:
 - Web server role
 - Windows internal database
 - Microsoft .NET framework.

SharePoint Server 2007 ceases to run if you uninstall these components.

- The account that you select for installing Office SharePoint Server 2007 must be a member of the Administrators group on every server where installed. You can, however, remove this account from the Administrators group on the servers after installation.
- You must install SharePoint Server 2007 on the same drive on all load-balanced front-end Web servers
- To deploy SharePoint Server 2007 in a server farm environment, you must provide credentials for several different accounts.
- All the instances of Office SharePoint Server 2007 in the farm must be in the same language.

- You must use the **Complete** installation option on all computers you want to be index servers, query servers or servers that run Excel calculation services. This option can also be used for all servers in the farm. As a general rule, all servers within the farm with SharePoint installations (all servers except SQL server) can use the complete option. As an example, this allows for adding query and other roles to Web servers without reinstalling the entire program.
- If you want to have more than one index server in a farm, you must use a different Shared Services Provider (SSP) for each index server.

Server Requirements

The application and database tiers sections explain the roles of storage that apply to include database requirements. For more information on storage requirements for each type of database, see the “Storage Sizing and Design” section of this document.

Web Services Tier

Table 4 describes recommended requirements for installing SharePoint front-end Web servers.

Table 4. Front-end Web Server Requirements

<i>Component</i>	<i>Configuration</i>
Processor	Dual processors, each running at 3GHz or faster
RAM	2GB or more
Network	1Gbps network interface card (NIC)

Application Tier

SharePoint application servers provide shared services such as search queries, indexing content and Excel services. These servers provide these services to the entire farm environment. This tier also hosts SharePoint central administration and a number of other mandatory and optional service provider functions. This tier has specific storage requirements including index storage and search databases. These requirements are in addition to the standard configuration and content databases SharePoint uses in SQL. Also, each Shared Service Provider requires its own database for service-specific data in addition to a separate database for search-related data.

Application server roles include:

- Query
- Index
- Windows SharePoint services search
- Shared Service Provider (SSP)
- Excel calculation services
- Office SharePoint search

Table 5 describes the hardware requirements for application servers.

Table 5. Application Server Requirements

<i>Component</i>	<i>Configuration</i>
Processor	Dual processors, each running at 2.5 GHz or faster
RAM	4GB or more
Network	1Gbps NIC

Database Tier

SQL Server provides database management for all SharePoint content and configuration data. The solution described in this white paper uses a SAN environment and Hitachi Adaptable Modular Storage 2000 family storage system mapped to the SQL 2008 server. Table 6 describes minimum hardware recommendations for a SQL 2008 implementation.

Table 6. Database Server Requirements

<i>Component</i>	<i>Configuration</i>
Processor	AMD Opteron, AMD Athlon 64, Intel Xeon with Intel EM64T support with Intel Pentium IV with EM64T support running at 2GHz or faster
RAM	2GB or more
Network	1Gbps NIC

SQL 2008 Installation

For SharePoint, it's most beneficial to make improvements in the following areas of SQL Server 2008:

- **Backup performance** — By implementing database backup compression.
- **Security** — By implementing Transparent Data Encryption (TDE).
- **Resource management** — By using SQL Server 2008 Resource Governor.

For more information, see the Microsoft TechNet articles, "[Plan for and design database storage and management](#)" and "[Integration of SQL Server 2008 and Office SharePoint Server 2007](#)."

This section provides key considerations for installing a single instance SQL Server 2008 application to prepare for SharePoint. For more information, see the Microsoft TechNet article, "[Quick-Start Installation of SQL Server 2008](#)."

Before installing SQL Server 2008, keep the following points in mind:

- Ensure that the prerequisite software is installed. Prerequisite software includes .NET Framework 3.5 and Installer 4.5. If this software is not installed, you can obtain it from Microsoft download centers.
- You can authenticate to SQL in one of two ways:
 - **Windows authentication mode** — Disables the SQL system administrator (sa) account
 - **Mixed authentication mode** — Enforces strict security requirements on SQL sa account as well as windows domain accounts

Select **Mixed Mode Authentication** for your instance of SQL Server if you want to use both Windows and built in SQL sa accounts to access the SQL instance, especially if your database administrators require the

use of the built-in SQL sa account. SharePoint requires the use of Windows Active Directory users to also have specific permissions to access and manage some of the databases required by WSS 3.0.

- During SQL installation, you must choose installation directories for database, transaction log and tempdb storage. Volume mount points or mounted drives are the suggested practice for use in directories when configuring SQL for SharePoint for two reasons. First, drive letter naming convention of limitation to 26 letters does not apply and second, mounted drives are better protected from changes in the event physical storage devices are added or removed. The database planning section provides examples of the storage LUs created for the different types of databases.

For more information, see the Microsoft TechNet article "[Volume Mount Points vs. Lettered Drives.](#)" Volume mount points are attached to folder names on NTFS volumes and are protected against system changes with normal drive letter configuration alterations.

- Complete the database and analysis services configuration steps for installation.

SharePoint Installation

You must install and configure SharePoint Server 2007 on all of the farm servers before you configure any SharePoint related services and create sites. You also must have SQL Server running on at least one database server before you install SharePoint on your Web and application servers. After all servers are configured, the SharePoint Product and Technologies configuration wizard completes the farm implementation and ties all tiers together. This includes running the wizard when adding any new front-end Web servers, index, search or any other servers to the farm.

Follow these high-level steps to deploy a farm of servers:

1. Ensure that the appropriate Windows updates applied and necessary prerequisite software are installed on all servers.

Note: This procedure assumes that SQL Server 2008 on the database server is already installed for the database tier.

Using the SharePoint 2007 installation media or a common network share, run setup on each application and web server in the farm.

Choose the **Complete** option for each installation. Note: The SharePoint Products and Technologies configuration wizard automatically runs upon completion of the SharePoint installation to finalize the integration of all tiers into the SharePoint farm.

The SharePoint central administration Web site launches upon of the wizard.

2. When SharePoint 2007 installations are complete, recheck Windows Updates to gather any pending patches or hot fixes that apply.
3. Complete all of the administrative tasks as shown on the central administration home page.

Instructions are provided for each step. After each step is complete, edit the task and provide completion details, such as who completed the activity, date of completion or deferral, and so on, in case you cannot complete the tasks at one time.

Front-End Network Load Balancing

When deploying multiple front-end Web servers, proper balance of the serving load is accomplished through a network load balancing (NLB) component. Your network load balancer, whether physical or virtual, must be configured prior to completing SharePoint installation. Windows 2008 Network Load Balancer must be manually installed on your Windows server because it is not installed by default. Load balancing on front-end Web servers ensures equal balancing across replicated Web servers to provide identical functionality for serving Web content.

For more information, see the Microsoft TechNet article "[Implementing a New Network Load Balancing Cluster.](#)"

Application Server Tier

Setting up this tier includes implementing all application structures of SharePoint including Shared Services Providers, central administration and Excel services in addition to any provisions for queries and search.

A single Shared Services Provider (SSP) is a logical grouping of shared services and their supporting resources. In Office SharePoint Server 2007, the SSP enables sharing services across multiple server farms, Web applications, and site collections. At least one SSP is required within a SharePoint 2007 farm. For more information, see the Microsoft TechNet article "[Create and configure Shared Services Providers.](#)"

Trace Log Configuration

The trace log can be useful for analyzing problems that might occur. For example, you can use events that are written to the trace log to identify configuration changes made in Office SharePoint Server 2007 before a problem occurred.

By default, Office SharePoint Server 2007 saves two days of events in the trace log files. This means that trace log files that contain events that are older than two days are deleted. Whether you are using the Office SharePoint Server Search service or the Windows SharePoint Services Search service, Hitachi Data Systems recommends that you configure the trace log to save seven days of events.

You can use the diagnostic logging page in central administration to configure the maximum number of trace log files to maintain and how long (in minutes) to capture events to each log file. By default, 96 log files are kept, each one containing 30 minutes of events:

96 log files * 30 minutes of events per file = 2,880 minutes or two days of events

You can also specify the location where the log files are written or accept the default path.

Validation

Like any other type of SQL Server deployment, it's important to test your environment to ensure that the current configuration achieves appropriate performance levels.

Microsoft offers the SQLIO and the SQLIOSim tools to help validate your SQL Server environment. Nevertheless, it is important that you test your own environment in a close to production mode. While SQLIO and SQLIOSim can provide some a generic view of the environment's capabilities, your specific database workload and user load might greatly differ from SQLIO. The Iometer tool is also a good option when you have deep knowledge of your specific database workload. For more information on utilizing tools to validate you SQL Server environment, see Microsoft's "[Deploying SQL Server 2005 with SAN #3](#)" article.

Always start by deploying the a test or development environment and appropriately test the SQL Server instance to ensure the correct levels of IOPs are achieved within acceptable latency ranges. It's important to establish baseline metrics for your environment so that you can monitor its performance and identify areas of concern.

SQLIOSim

The SQLIOSim utility tests data integrity capabilities of I/O paths. It simulates read patterns, write patterns, and the problem identification techniques of SQL Server. To perform these tasks, the SQLIOSim utility simulates a SQL Server system's user and system activity. The SQLIOSim utility performs this simulation independent of the SQL Server engine.

The utility generates log files that show execution and error details.

NOTE: Be careful to not specify the actual SQL Server database files for testing if your SharePoint databases are already prepared because this utility overwrites data with random test patterns, destroying current data.

For more information, see the Microsoft article "[How to use the SQLIOSim utility to simulate SQL Server activity on a disk subsystem.](#)"

SQLIO

SQLIO is a free Microsoft utility that measures I/O capacity performance. This tool has nothing to do with SQL Server but is valuable in stress testing and analyzing your SAN. Use this test in parallel among all application and SQL servers that touch your SAN within your SharePoint farm to determine if multiple servers running on the same SAN affect each other's performance.

You can download SQLIO at Microsoft's [Download Center](#). For more information, see the SQLServerPedia article "[SAN Performance Tuning](#)."

Iometer

Iometer is an open source performance analysis tool that includes a graphical interface that allows you to configure the workload, set parameters and start and stop tests to test disk performance. It collects the test results and produces the reports for analysis. Iometer measures performance under controlled load, as a workload generator and a measurement tool. It can be configured to emulate the disk or network I/O load of any program or benchmark, or can be used to generate entirely synthetic I/O loads. It can generate and measure loads on single or networked systems. Iometer can be used for measurement and characterization of the following metrics:

- Performance of disk and network controllers
- Bandwidth and latency capabilities of buses
- Network throughput to attached drives
- Shared bus performance
- System-level hard drive performance
- System-level network performance

To download Iometer, visit [SourceForge.net](#).

SharePoint 2007 Test Data Population Tool

This open source capacity planning and performance testing tool populates data for testing SharePoint deployments. You can also download and create test scripts that can be used for additional and enhanced performance and load testing while viewing results either from a monitor such as Iometer or the Visual Studio Team Explorer.

To download this tool, visit [CodePlex.com](#).

Visual Studio Team System Test Edition

Visual Studio Team System Test Edition allows you to create and view results of testing scripts that can show you're your SharePoint environment is operating. The main purpose of this testing environment is to combine tests such as the data population tool and built-in test scripts to evaluate the SharePoint farm's operational conditions. For more information, visit Microsoft's [Visual Studio Team System Test Edition Web page](#).

Confirming SharePoint 2007 Functionality

It's important to confirm that your SharePoint 2007 environment is functioning properly before making it live. Test the functionality of your SharePoint farm environment by creating and manipulating content using the existing configurations of storage and roles. If you can create and manipulate content without errors, you can rest assured that your SharePoint environment is functioning correctly.

Validate Site and Content Creation

To validate site and content creation, follow these steps:

1. Choose **Start > Administrative tools** in SharePoint central administration.
2. Click the **Application Management** tab.
3. Test creation of a Web application.
4. Test creation of a site collection.
5. Test deleting the same application and collection.

Validate Sub-site, Library and List Creation

To validate creation of sub-sites, testing libraries and list functionality, follow these steps:

1. Click the **Shared Documents** link to upload, download, delete or edit files within a library.
2. Click the **Calendar** link and add events or other calendar items.
3. Tasks are similar as each of these lists include an **Add** or **Edit** link to provide new data or change existing content.
4. Click the **Sites** link and add sub-sites to the top level site, alter or add permissions for any site, list or library,

Hitachi Storage Management Software

Hitachi Data Systems offers software to provide performance and specification recommendations and validations. This section addresses the value in determining your environment's characteristics and throughput.

Hitachi Storage Navigator Modular 2 Software

Hitachi Storage Navigator Modular 2 software is the integrated interface for Adaptable Modular Storage 2000 family firmware and software features. Use it to take advantage of all of the 2000 family's features. Storage Navigator Modular 2 software provides both a Web-accessible graphical management interface and a CLI to allow ease of storage management.

Storage Navigator Modular 2 software is used to map security levels for SAN ports and virtual ports and for inter-system path mapping. It is used for RAID-level configurations, for LU creation and expansion and for online volume migrations. It also configures and manages Hitachi replication products. It enables online microcode updates and other system maintenance functions and contains tools for simple network management protocol (SNMP) integration with enterprise management systems. SNMP is a network protocol that is used to [monitor](#) for conditions that require administrative attention.

Hitachi Performance Monitor feature provides detailed, in-depth storage performance monitoring and reporting of Hitachi storage systems including drives, logical volumes, processors, cache, ports and other resources. It helps organizations ensure that they achieve and maintain their service level objectives for performance and availability, while maximizing the utilization of their storage assets. Performance Monitor feature's in-depth troubleshooting and analysis reduce the time required to resolve storage performance problems. It is an essential tool for planning and analysis of storage resource requirements.



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 2009. All Rights Reserved.
AS-018-00 October 2009