HITACHI
Inspire the Next

Microsoft
GOLD CERTIFIED
Partner

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

Best Practices and Reference Architecture Guide

By Art LaMountain

October 2009

Hitachi Storage Solutions

Hitachi Data Systems

Executive Summary

Deploying a SharePoint 2007 environment can be a costly and complex endeavor without proper planning and deployment guidance. However, when effectively implemented, the benefits of using the Hitachi Adaptable Modular Storage 2000 family are significant, and include reduced storage and server footprints, reduced power costs and simplified database infrastructure management.

Businesses of all sizes need storage solutions for enterprise applications such as SharePoint that are easy to deploy and maintain, are highly available, provide flexible scalability, deliver predictable performance, and introduce consolidation-related savings. This paper is written for storage and database administrators in large and enterprise organizations who use SharePoint 2007 for content management and collaboration services and those interested in the benefits that the Hitachi Adaptable Modular Storage 2000 family offers for storage. It assumes working knowledge of SharePoint 2007, security technologies like Windows Active Directory, Kerberos and domain name services and basic understanding of configuration and architecture of the 2000 family.

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

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

Best Practices and Reference Architecture Guide

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

This solution describes a building block architecture that uses the Hitachi Adaptable Modular Storage 2000 family to provide these benefits:

- A robust, scalable architecture for storing, managing and accessing large – and growing – amounts of business information
- Consolidated files and storage to ease management burden and reduce operational overhead
- Increased productivity through improved collaboration, high availability, and easier access to archived data

This white paper provides best practices and key considerations for planning, configuring and deploying SharePoint with a 2000 family storage system using a building block approach that simplifies scalability and management of the IT infrastructure. It addresses planning considerations including sizing, platform hardware foundations, high availability requirements and security.

This document is written for storage and database administrators in large and enterprise organizations who use SharePoint 2007 for content management and collaboration services and those interested in the benefits that the Hitachi Adaptable Modular Storage 2000 family offers for storage. It assumes working knowledge of SharePoint 2007, security technologies like Windows Active Directory, Kerberos and domain name services, and basic understanding of configuration and architecture of the 2000 family.

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.

The Hitachi Adaptable Modular Storage 2000 family is the only midrange storage product with symmetric active-active controllers that provide integrated, automated hardware-based front-to-back-end I/O load balancing. Both controllers in an 2000 family storage system are able to dynamically and automatically assign the access paths from the back of the controller to the logical unit (LU). All LUs are accessible regardless of the physical port or the server from which the access is requested. Utilization rates of each controller are monitored so that a more even distribution of workload between the two controllers can be maintained.No other midrange storage product that scales beyond 100TB has a serial attached SCSI (SAS) drive interface. The new point-to-point back-end design virtually eliminates I/O transfer delays and contention associated with Fibre Channel arbitration and provides significantly higher bandwidth and I/O concurrency. It also isolates any component failures that might occur on back-end I/O paths.

SharePoint 2007

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

The size of your environment depends largely on the quantity of users and complexities of the applications required. SharePoint 2007 deployments can range from a standalone environment with a single server instance for fewer users to a farm environment that can grow up to many thousands of users with needs for more collaboration and functionality. A farm allows scaling by adding servers to the environment as the number of users or amount of content increases. Environmental additions such as servers within each tier can be made at any time simply by connecting the additional hardware and running the SharePoint configuration wizard.

Figure 1 illustrates the foundation environment consisting of three tiers plus surrounding Windows infrastructure.

Figure 1. SharePoint Tiers and Windows Infratsructure

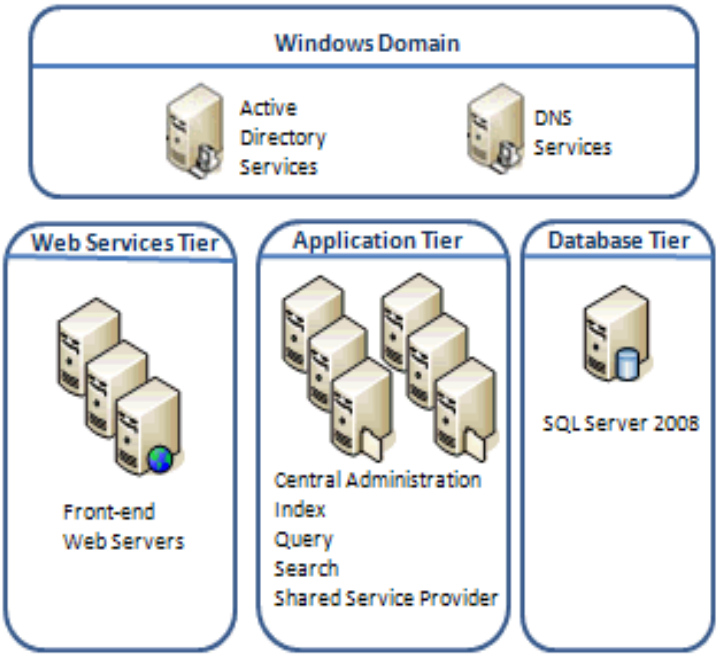

Each tier can be expanded by adding the roles shown. This can be achieved by adding SQL servers, adding Web servers to provide better throughput and network load balancing to end users, or adding application roles to integrate new functions within the power of a unique server.

Storage Reference Architecture

This section describes important considerations to keep in mind when planning and implementing this solution using the Hitachi Adaptable Modular Storage 2000 family.

This solution uses a 1.5TB reference architecture that stores all SharePoint databases and index functionalities with corresponding transaction log and `tempdb` data. It is important to deploy your solution in a test environment to ensure that you meet service level agreements (SLAs). This might require you to increase the amount of spindles.

The 2000 family includes features that make it easy to configure and implement a SQL 2008 environment. The placement of RAID groups is no longer dependent on a back-end loop structure. The switched, 3Gb/sec SAS back-end architecture in the 2000 family also allows bandwidth to all disks in a tray by utilizing 16 wide links. The switched architecture provides fault tolerance so that if one of the wide links fails, the SAS controller can detect the failure and continue to use the remaining links. The ability to use either SAS or SATA drive technology in the 2000 family expansion tray offers additional flexibility.

Tiers and Database Connectivity

Hitachi Data Systems testing used 15 100GB content databases for a medium-sized SharePoint deployment.

Properly configuring the Hitachi Adaptable Modular Storage 2000 family for scalability and performance is important to meet your users' expectations and organizational service level agreements. Figure 2 shows the SharePoint components known as tiers and additional interfacing hardware for the 2000 family.

Figure 2. Server Roles for Fibre Channel and Network Connectivity

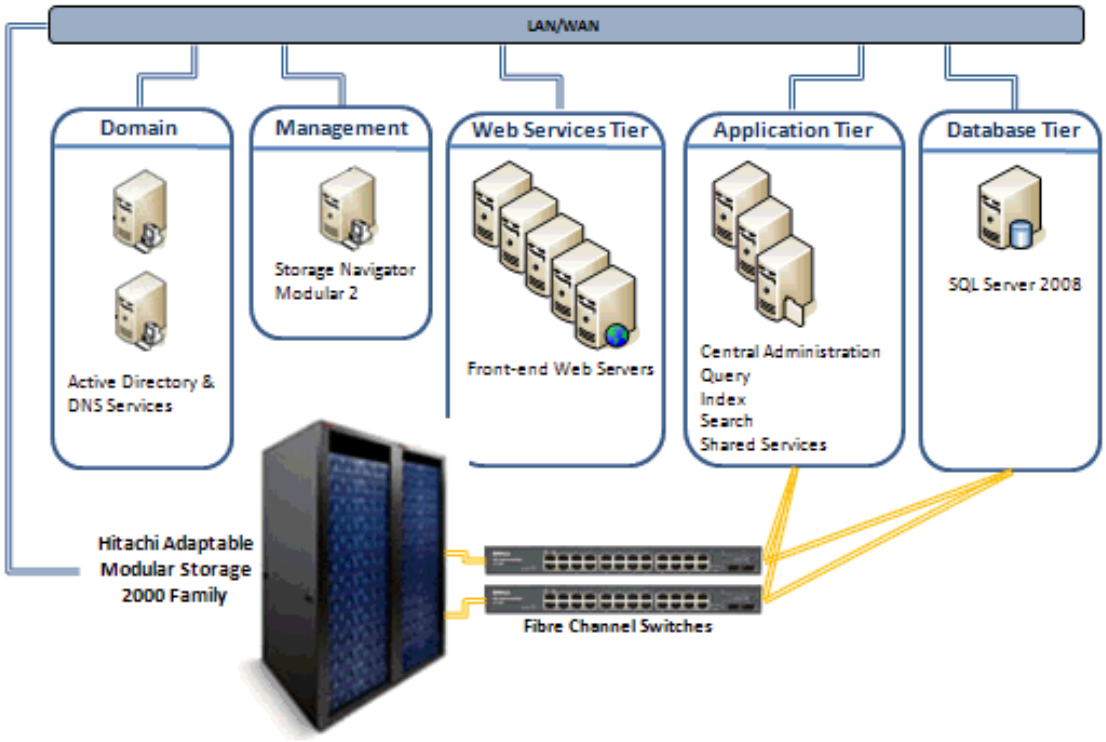

It is important to understand architectures of both SharePoint and the external SAN storage used for SQL Server. Indexes are maintained an NTFS partition. For more information, see the *Implementing SharePoint 2007 with the Hitachi Adaptable Modular Storage 2000 Family Planning and Deployment Guide* white paper.

The deployment tested for this solution used 15 SAS 450GB 15K RPM drives per tray with a single spare. Multiple content databases, the search database, index and their associated LUs were deployed on two dedicated RAID-1+0 (2D+2D) groups. A third RAID-1+0 (2D+2D) group was dedicated for additional content. Each of these groups contained enough storage (approximately 20 percent) for future growth of existing databases. The transaction logs for this deployment were hosted on a single RAID-1 (mirrored) (1D+1D) groups. Because of the SAS back-end architecture, these RAID groups can technically be placed anywhere on the storage system. To make configuration simple, all of the RAID groups were configured allowing Storage Navigator Modular 2 software to automatically assign the disks.

In general, the goal is to balance the load across the ports and controllers. In addition, the port assignment for the alternate path for each LU assumes the use of Hitachi Dynamic Load Manager software for automatic failover in the event of a failure along the primary path. Figure 3 shows the single tray relationship of RAID groups to database and log type locations.

Figure 3. Hitachi Adaptable Modular Storage 2000 Family with SAS 450GB Drives

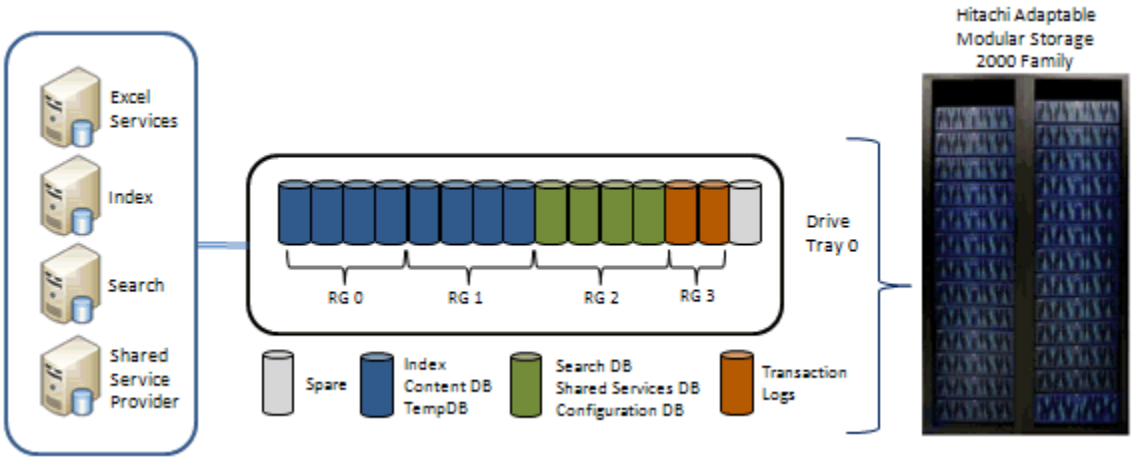

Table 1 describes the RAID groups created to sustain the capacity required for the deployment. Staging and production environments should closely resemble each other for proper testing and deployment. The development environment does not necessarily require the same performance and failover characteristics as a production implementation.

Table 1. RAID Groups for Database Capacity Deployment

RAID Group	*RAID Level*	*Configuration*	*Number of Drives*	*Use*
0	1+0	2+2	4	Content databases and index and `tempdb`
1	1+0	2+2	4	Content databases
2	1+0	2+2	4	Shared services databases, search database and configuration database
3	1	1+1	2	Transaction logs

The configuration used for the deployment described in this paper assumes eight paths between the storage system and the database server distributed equally among each database type. Note that `tempdb` and transaction log files utilize a unique primary physical port from the actual database file. This ensures balancing so that the physical database has its own primary path prioritization. Table 2 shows typical LU configuration and distribution of port assignments to equalize the load between storage and the database server. This table shows the first four of a possible total of 15 content databases. Additional databases maintain assignments in the same RAID structure beginning with the next sequential numbering schema for LUs and RAID groups.

Table 2. RAID Groups for Database Capacity Deployment

LU Assignment	*LUN*	*RAID Group*	*LU Size (GB)*	*Port Assignments*	
tempdb	0	0	10% of all database sizes on this instance	0A	1A
Database 0					
Content DB0	1	0	100	1A	0A
Transaction log	2	3	20	0A	1A
Database 1					
Content DB1	3	0	100	0B	1B
Transaction log	4	3	20	1B	0B
Database 2					
Content DB2	5	0	100	1C	0C
Transaction log	6	3	20	0C	1C
Database 3					
Content DB3	7	0	100	0D	1D
Transaction log	8	3	20	1D	0D
Index	9	0	400	1A	0A
SSP database					
Search database	9	2	400	1A	0A
Transaction Log	10	3	80	0A	1A
Configuration database					
Configuration database	13	2	5	0B	1B
Transaction Log	14	3	1	1B	0B

Scalability

Scaling your SharePoint environment includes adding components and configurations to numerous areas including adding font-end servers for performance, adding site collections that require additional content databases, implementing additional applications and more. This paper focuses on scalability issues that require adding storage databases to the 2000 family storage system. The deployment tested for this solution uses two RAID groups for 15 100GB content databases, and a third RAID group for the first SSP and search databases.

When your SharePoint environment requires additional storage, whether it is needed for content or additional search indexing, provide added RAID groups as outlined in the Storage Refernce Architecture section. Figure 4 illustrates additional storage for up to 15 content databases and another Shared Service Provider with search. The configuraton includes the addition of a three new RAID-1+0 (2+2) groups and corresponding LUs as RAID groups 4 through 7, with RAID group 7 set aside for additional transaction logs.

Figure 4. Adding Storage for Content Growth

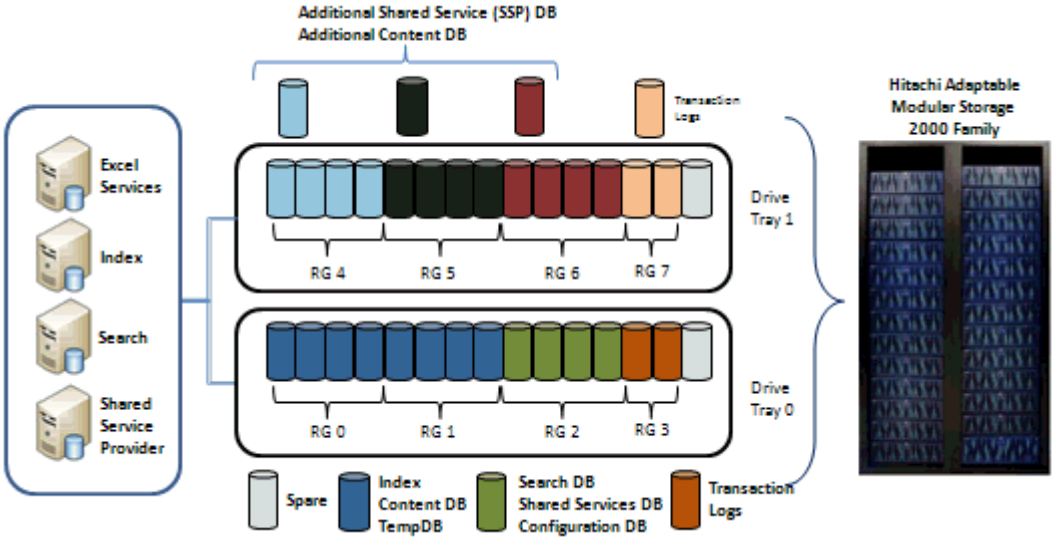

Best Practices

This section describes best practices that ensure a secure, high-performance and scalable SharePoint deployment.

Key Considerations

Follow these recommendations for the management and configuration of the SAN environment in relationship to Fibre Channel switch zoning and HBA configurations:.

- Use at least two HBAs and place them on different buses within the server to distribute the workload over the server's PCI bus architecture.
- Use of dual SAN fabrics, multiple HBAs and host-based multipathing software when deploying business-critical SQL Server applications. Two or more paths from the SQL and application servers connecting to two independent SAN fabrics is essential for ensuring the redundancy required for critical applications.
- Zone your fabric appropriately for multi-path HBA to storage ports. Use at least two Fiber Channel switch fabrics to provide multiple independent paths to the 2000 family storage systems to prevent configuration errors from bringing down the entire SAN infrastructure.
- For large bandwidth requirements that surpass an HBA's port capability, use round robin as the load-balancing setting for either Hitachi Dynamic Link Manager software or Windows Server 2008 native MPIO.
- Maintain multiple copies of your database and log files with Hitachi ShadowImage® Hetereogenous In-System Replication software or Copy-on-Write Snapshot software using the Microsoft Virtual Device Interface (VDI) framework for quick recovery.

Database Design Considerations

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

Every SQL Server database has at least one log file that records database modifications made by each transaction. It is a critical component of the database for availability and data integrity. In the event of a system failure, the active transaction log, at a minimum, is required to bring the database back to a consistent state.

The transaction logs are written before the data records are updated to the database file via the checkpoint process. The logs can be used to roll back transactions if corruption is later discovered or to recover a committed transaction if an error writing to the database occurs.

Response time and performance are critical considerations for separating the transaction log from the database files. Microsoft suggests aiming for log I/O response times between 1 and 5 milliseconds. Hitachi's optimized caching and proper storage design ensure that the logs can be written without delay.

One of the features built into the 2000 family is the ability to optimize physical I/O based on recognition of I/O patterns. The 2000 family controller can make optimized timing decisions about when to move the data between mirrored, protected cache and physical disks when it encounters a series of I/O requests. By combining multiple logical I/O requests into a single physical I/O or by optimizing the order of individual reads and writes, the 2000 family can significantly increase overall performance.

SQL Server `tempdb` Files

SQL Server `tempdb` files are used for storage of temporary data structures and can be subject to intense and unpredictable I/O. Many best practice recommendations suggest locating `tempdb` files on separate RAID groups from the database and using the fastest disks available. This is generally a safe recommendation because the load on tempdb is highly dependent upon database and application design. However, if the `tempdb` load is well understood and monitored regularly, testing shows that it can reside on the same RAID group as the primary database without adverse effects. Accordingly, if the environment does not have sufficient physical disk I/O resources to meet the combined requirements of `tempdb` and the database files, performance for all databases in the SQL Server instance degrades. Therefore, you need a good understanding of your `tempdb` usage, regardless of where you choose to place the `tempdb` files. The 1.5TB environment used for this paper does not require separation of the `tempdb` files. However, if you establish higher performance requirements, you can create an additional RAID group (2+2) for the `tempdb` data.

Database Performance Considerations

Performance considerations include monitoring and measuring Average Disk sec/Read and Average Disk sec/Write**.** The average response time in milliseconds for a logical or physical disk can be affected by I/O size, RAID configuration and other factors in the data path.

Lower measures of disk latency are better but can vary dependent on the size and nature of the I/Os being issued. These numbers also vary across different storage configurations. Cache size and utilization can greatly affect this.

SQL Server database file I/O is composed of random small record reads and writes. A database might include only a single database file, while those designed to support heavy transactional workloads or large schemas might use a variety of filegroup architectures to improve performance, operational convenience or availability. Acceptable database I/O response times typically range between 5 and 20 milliseconds.

For more information, see the *Hitachi Adaptable Modular Storage 2000 Family Best Practices with Microsoft® SQL Server for Online Transaction Processing Applications* white paper.

Deployment Considerations

Keep these design considerations for sizing databases for SharePoint environment in mind:

- **Database log files** — Disk space for log files varies based on log settings and the number of databases. For more information, see the Microsoft TechNet article "Physical Database Storage Design."
- **Configuration database** — The configuration database generally does not grow past 1.5GB in size. This is an estimated maximum size, not a hard limit.

- **Content databases** — Estimate the initial volume of content that will be stored in content databases. Consider the following factors:
  - Multiply the value of the size of initial content by 1.2 to get the value for the size of stored content in a SQL database.
  - If versioning is used for documents, a copy of each version is stored in the database.
  - Microsoft's recommended limit is 100GB.
- **Future growth** — Plan for twice the amount of data that you initially plan to experience. Allow at least 25 percent free space for each hard disk or volume.

For more information about how to estimate the size of subsequent content, search and SSP databases, see the following documents:

- Hitachi Data Systems *Implementing SharePoint 2007 with the Hitachi Adaptable Modular Storage 2000 Family Planning and Deployment Guide* white paper
- Microsoft TechNet articles:
  - "Best Practices for Capacity Management"
  - "Best Practices Operational Excellence"
  - "Best Practices Resource Center"
  - "Best Practices for Search"
  - "Estimate performance and capacity requirements for SharePoint"

Database Configuration Best Practices

To ensure that your databases are prepared for scalability and proper health, follow these best practices:

Common to all databases:

- Use RAID-1 (1D+1D) for transaction logs and RAID-1+0 (2D+2D) groups for index, search, configuration and content databases.
- Create one `tempdb` file per CPU core (see the following note) and make all files equal in size.

  Note: Count dual-core processors as two CPUs for this purpose. Create files only in the primary file group for the database. The host can write to each `tempdb` file at the same time due to the multi-core capability. This prevents `tempdb` from becoming a bottleneck. For more information, see the SQL Server Developer Center article "*Optimizing `tempdb` Performance*."
- Pre-size data and log files rather than relying on autogrowth. Enabling autogrowth to meet deployment needs for safety purposes is fine as long as you manage the growth of the data files with proper planning. Additionally, turn on the `MAXSIZE` setting for each of the DB files in order to ensure none of them grow to a point where it uses the entire space available on a disk.

Configuration database:

- Do not overprovision the configuration database with storage. The configuration database is fairly static after it is created and should not grow past 1.5GB.

Content database:

- Use multiple data files for content and SSP search databases. (See the note in the previous list.)
- Create data files of equal size.
- Provide an autogrowth value of a fixed number of megabytes instead of a percentage for databases exceeding 100GB. This reduces the how often these file sizes must be increased. When you are planning content databases smaller than the recommended size (100GB), set the databases to 100GB when they are created by using the `MAXSIZE` setting.

- If disk space is limited or databases cannot be sized, configure the autogrowth value to a fixed percentage such as 10 percent for databases under 500GB and to a fixed number of megabytes if a database exceeds 500GB.
- To enhance the performance of SharePoint 2007, Microsoft discourages the use of content databases larger than 100GB. For databases larger than 100GB, follow these recommendations:
  - Use multiple site collections for the data, each owning owning a unique content database.
  - Whenever possible, Microsoft recommends splitting content from a site collection that is approaching 100GB into a new site collection in a separate content database to avoid performance or manageability issues.

Microsoft provides specific recommendations for sizing and organizing SharePoint content databases, including these:

- Limit content databases to 100GB.
- Use site collection quota management to monitor and prevent excess storage requirements.

For more information, see the Microsoft TechNet article "Best Practices for capacity management."

Database Maintenance Best Practices

To ensure that your databases are properly maintained, follow these best practices:

- Monitor the database server to make sure that it is responding appropriately and is not overloaded. Key performance counters to monitor include the following:
  - **Network wait queue** — 0 or 1
  - **Average disk queue length (latency)** — Less than 20ms
  - **Memory used** — Less than 70 percent
  - **Free disk space** — More than 25 percent for content growth
- Do not auto-shrink databases or set up any maintenance plans that programmatically shrink your databases.
- Consider shrinking a database only when 50 percent or more of the content in it has been removed by user or administrator deletions. Shrinking databases is very resource intensive and requires careful scheduling.
- Consider database shrinking only for content databases. The configuration, central administration, SSP and search databases do not usually experience enough deletions to contain sufficient free space.
- Avoid shrink requirements by including growth allocations in your capacity planning, including an overhead allocation of 10 to 20 percent.

Management Best Practices

As the amount of content or numbe of users increase, it is important to monitor and manage your enivornment to ensure that databases and performance changes do not negatively affect operations. For example, you might increase the amount of site collections for company acquisitions or internal departmental growth, requiring more storage space and generating heavier user access. After monitoring IOPS and IIS characxteristics, you might determine that additional storage LUs are needed or additional front-end Web servers might help balance user traffic within the SharePoint farm.

Both Microsoft and Hifachi offer numerous software management products that can assist in the management of your environment from performance to configuration.

Hitachi Storage Management Software

Hitachi offers several tools for managing and monitoring the performance of your deployment in general and your 2000 family storage in particular:

- **Hitachi Tuning Manager software —** A real-time software monitor that you can use to check the current state of all the host, file system, database, SAN and storage resources that are being used by applications. It lets you contrast data about the monitored resources with a historical view of the normal, baseline performance of those resources that was previously stored in the database. The ability to query a historical database for performance and capacity trend analysis for each component of a SAN enables you to easily correlate the current changes in performance with recent changes to the physical configuration, software, or workload or with other environmental changes that might be causing changes in an application's performance.
- **Hitachi Storage Navigator Modular 2 software —** Provides an integrated interface for the Hitachi 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 command-line interface to simplify storage management.

  Use Storage Navigator Modular 2 software to map security levels for SAN ports and virtual ports and for inter-system path mapping. You can also use it for RAID-level configurations, LU creation and expansion and online volume migration. It also configures and managers Hitachi Replication products,enables online microcode updates and other system maintenance functions and contains tools for Storage Navigator Modular 2 software integration with enterprise management systems.

  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. 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.
- **Hitachi Dynamic Link Manager software** — Manages paths between a host and a storage subsystem. Dynamic Link Manager software distributes loads across multiple paths and switches a load to another path if a failure occurs in the path that is currently being used, thus improving system reliability. Dynamic Link Manager software manages the access paths to these storage subsystems.
- **Hitachi QoS for Microsoft® SQL Server** — Identifies in real time how storage infrastructure performance, capacity, and availability impact SQL Server. It auto-discovers your SQL Server application infrastructure, charting the dependencies of SQL Server instances and databases on hosts, host bus adaptors (HBAs), SAN switches, and disk systems. By drilling down from SQL Server performance and availability issues to the underlying cause and issuing alerts when storage infrastructure changes affect SQL Server, QoS for SQL Server helps administrators satisfy service-level requirements, respond more efficiently to help-desk calls, increase application uptime, and plan for effective disaster recovery.
- **Hitachi Device Manager software** — Provides a single platform for centrally managing, configuring and monitoring Hitachi storage systems. By significantly boosting the volume of storage that each administrator can manage, the single-point-of-control design of Device Manager software can help raise storage management efficiency in these environments as well as reduce costs. Easy-to-use Device Manager software views storage resources logically, while maintaining independent physical management capabilities. By offering a continuously available view of actual storage usage and configuration, Device Manager software allows administrators to precisely control all managed storage systems.

Microsoft Management Software

Microsoft offers the following tools to help you monitor and manage the health of your SharePoint deployment:

- **MOSS 2007 Best Practice Analyzer for Windows SharePoint Services 3.0** — This software collects information from your SharePoint environment, including various settings and other values from SQL, registry and performance monitor, to create detailed reports to help administrators review and enhance performance, scalability and uptime.
- **System Center Operations Manager (SCOM) 2007** — Manages and monitors a wide variety of services and health indicicators of the overall IT environment, affording IT administrators the ability to proactively react to conditions that might impact performance or availability of potentially serious conditions. For more information, see the Microsoft TechNet article **"**System Center Operations Manager 2007."

Hitachi Data Systems Corporation

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

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

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

AS-019-00 October 2009