

# Deploying Microsoft® SharePoint® 2010 on the Hitachi Virtual Storage Platform with Hitachi Dynamic Tiering

Reference Architecture Guide

*By Mark Crowley*

*July 2011*



## Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to [SolutionLab@hds.com](mailto:SolutionLab@hds.com). Be sure to include the title of this white paper in your email message.

# Table of Contents

<b>Solution Overview .....</b>	<b>4</b>
<b>Key Solution Components.....</b>	<b>6</b>
Hitachi Virtual Storage Platform .....	6
Hitachi Dynamic Provisioning.....	6
Hitachi Dynamic Tiering .....	7
Hitachi Compute Blade 2000.....	7
Hitachi Dynamic Link Manager Advanced .....	7
Microsoft SharePoint Server 2010 .....	8
Microsoft SQL Server 2008 R2 .....	8
Brocade Storage Area Network Switches .....	8
<b>Solution Design .....</b>	<b>8</b>
High-level Infrastructure .....	8
Storage Architecture.....	10
SharePoint Architecture .....	15
Server Architecture.....	18
SAN Architecture .....	22
Physical and Virtual Network Architecture .....	23
<b>Engineering Validation.....</b>	<b>25</b>
Test Methodology.....	25
Test Results.....	26
<b>Conclusion .....</b>	<b>27</b>

# Deploying Microsoft® SharePoint® 2010 on the Hitachi Virtual Storage Platform with Hitachi Dynamic Tiering

## Reference Architecture Guide

Businesses increasingly rely on Microsoft SharePoint Server 2010 as a collaboration tool to enable employees to work together, share information and manage content across countries, continents and time zones. Planning and deploying SharePoint in enterprise environments is a complex endeavor requiring many hours of effort. A robust deployment requires complex sizing and location determinations. These have an immediate and profound impact on overall system performance.

In a Microsoft SharePoint 2010 environment, systems administrators constantly deal with performance and sizing challenges. These relate to the multitude of documents and indexed data that is typically referred to as “SharePoint sprawl.”

Simplify these performance and management challenges with the Hitachi Virtual Storage Platform, an enterprise-class storage system. This platform features Hitachi Dynamic Provisioning and Hitachi Dynamic Tiering, which supports thin provisioning, wide striping, and automatically tiered storage functionality.

The solution in this reference architecture plans and deploys a 200,000 user Microsoft SharePoint 2010 farm with 20 site collections using 20 content databases. The site collections are hosted on a single web application. The use of the latest technologies from Microsoft and Hitachi Data Systems effectively uses the resources available in your data center while maintaining the high availability and performance levels required for a SharePoint environment.

The intended use of this reference architecture is by IT personnel responsible for planning and deploying Microsoft SharePoint infrastructure and associated storage. It assumes some familiarity with storage area networks, Microsoft Windows 2008 R2, Microsoft Hyper-V, Microsoft SharePoint 2010, and the features found on the Hitachi Virtual Storage Platform.

## Solution Overview

This is a tested reference architecture that provides a cost effective storage configuration for Microsoft SharePoint 2010 data on a Hitachi Virtual Storage Platform.

Hitachi Data Systems testing used four Hitachi Compute Blade 2000 server blades.

- Blade 0 and Blade 1 are Hitachi Compute Blade 2000 E57A1 blades, each with 16 processor cores and 128GB RAM. These blades are divided up into the following logical partitions:
  - **Blade 0, 4 cores, 32GB RAM**—This is used for the SharePoint 2010 Search Crawl role. Microsoft SQL Server 2008 R2 is installed for the search databases.
  - **Blade 0, 4 cores, 6GB RAM**—This is used for Microsoft Active Directory and DNS.
  - **Blade 1, 12 cores, 96GB RAM**—This is used for Microsoft SQL Server 2008 R2 which was the central SQL server for the SharePoint farm.
- Blade 6 and Blade 7 are Hitachi Compute Blade 2000 E55A2 blades, each with 12 processor cores and 96GB RAM. These are setup as Microsoft Windows 2008 R2 Hyper-V hosts.
  - The Hyper-V hosts each contain four Microsoft Windows 2008 R2 guest virtual machines (VM). All guest VMs run Microsoft SharePoint 2010, serving as the web front-end role and as the search query component (mirrors to the main query engine).
  - One guest VM on each host also hosted central administration.

The storage infrastructure uses the Hitachi Virtual Storage Platform to provide the following:

- Robust, scalable architecture hosted in SharePoint Server environments to store, manage, and access large and growing amounts of information
- Wide striping and thin provisioning using Hitachi Dynamic Provisioning to minimize storage footprint and to maximize performance
- Automatic data migration using Hitachi Dynamic Tiering to locations on the VSP where performance is appropriate for the application needs

Figure 1 shows the infrastructure used to host the Microsoft SharePoint test environment.




Figure 1

# Key Solution Components

The following are the key components used for the solution.

## Hitachi Virtual Storage Platform

The Hitachi Virtual Storage Platform is a 3D scaling storage platform. With the unique ability to scale up, scale out, and scale deep at the same time in a single storage system, the Virtual Storage Platform flexibly adapts for performance, capacity, connectivity, and virtualization.

- **Scale Up**—Increase performance, capacity, and connectivity by adding cache, processors, connections, and disks to the base system.
- **Scale Out**—Combine multiple chassis into a single logical system with shared resources.
- **Scale Deep**—Extend the advanced functions of the Virtual Storage Platform to external multivendor storage.

The switch matrix architecture of the Virtual Storage Platform makes all of this possible. It connects the basic components, front-end directors, back-end directors, global cache modules, and virtual storage directors. Adding redundant pairs of directors and cache modules can be done without any disruption to connected host servers. All these resources are coupled tightly through a global cache that creates a common pool of storage resources. These resources can include external storage that is connected through front-end director initiator ports.

For more information, see the [Hitachi Virtual Storage Platform](#) on the Hitachi Data Systems website.

## Hitachi Dynamic Provisioning

On the Hitachi Virtual Storage Platform, Hitachi Dynamic Provisioning provides wide striping and thin provisioning functionalities.

Using Hitachi Dynamic Provisioning is similar to using a host-based logical volume manager (LVM), but without incurring host processing overhead. It uses one or more wide-striping pools across many RAID groups on the Virtual Storage Platform. Each pool has one or more virtual volumes (DP-VOLs) of a user-specified logical size of up to 60TB created against it.

Deploying Hitachi Dynamic Provisioning avoids the routine issue of hot spots that occur on logical devices (LDEVs). These occur within individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. This distributes the host workload across many RAID groups, which provides a smoothing effect that dramatically reduces hot spots.

Hitachi Dynamic Provisioning has the benefit of thin provisioning. Physical space assignment from the pool to the DP-VOL happens as needed using 42MB pages. This continues up to the logical size specified for each DP-VOL. There can be a dynamic expansion or reduction of pool capacity without disruption or downtime. An expanded pool can be rebalanced across the current and newly added RAID groups for an even striping of the data and the workload.

For more information, see the [Hitachi Dynamic Provisioning datasheet](#) and [Hitachi Dynamic Provisioning](#) on the Hitachi Data Systems website.

## Hitachi Dynamic Tiering

Hitachi Dynamic Tiering eliminates manual data classification and movement between storage tiers. This optimizes tiered storage usage while improving performance.

Instead of manually provisioning space from several storage technologies with different performance and cost characteristics, Hitachi Dynamic Tiering enables the management of multiple storage tiers as a single entity. By leveraging the existing features of Hitachi Dynamic Provisioning, Hitachi Dynamic Tiering presents a virtual volume with embedded smart tiering. It monitors access and moves data at the 42MB page level.

Breaking the volume into pages, Hitachi Dynamic Tiering automatically moves infrequently referenced pages to lower cost tiers of storage. Moving pages instead of entire data sets or files reduces the time and storage space required to migrate data.

After an initial setup process, Hitachi Dynamic Tiering monitors data access in real time. It makes decisions on moving data between the available storage tiers based on actual use. Using this approach, Hitachi Dynamic Tiering improves the availability and performance of your storage systems and the applications using that storage.

Hitachi Dynamic Tiering on the Virtual Storage Platform allows a single pool to contain tiers made up of differently-arranged RAID groups using any type of disk. It manages data migration between the various tiers within a pool automatically. This eliminates most user management of storage tiers within a storage system, and maintains peak performance under dynamic conditions without storage administrator intervention.

For more information, see [Hitachi Dynamic Tiering](#) on the Hitachi Data Systems website.

## Hitachi Compute Blade 2000

The Hitachi Compute Blade 2000 is an enterprise-class blade server platform. It features the following:

- A balanced system architecture that eliminates bottlenecks in performance and throughput
- Embedded logical partition virtualization
- Configuration flexibility

Hitachi embeds logical partitioning virtualization in the firmware of the Hitachi Compute Blade 2000 server blades. This proven, mainframe-class technology combines Hitachi's logical partitioning expertise with Intel VT technologies to improve performance, reliability, and security. Embedded logical partition virtualization does not degrade application performance and does not require the purchase and installation of additional components.

## Hitachi Dynamic Link Manager Advanced

Hitachi Dynamic Link Manager Advanced is a software package that combines all the capabilities of Hitachi Dynamic Link Manager and Hitachi Global Link Manager into a comprehensive multipathing solution. Installed on each Microsoft SQL Server host, Dynamic Link Manager includes capabilities such as the following:

- Path failover and failback
- Automatic load balancing to provide higher data availability and accessibility.

Hitachi Dynamic Link Manager, used for storage area network multipathing, was configured with the round-robin multipathing policy. This round-robin load balancing algorithm automatically selects a path by rotating through all available paths. This balances the load across all available paths and optimizing IOPS and response time.

For more information, see [Hitachi Dynamic Link Manager](#) on the Hitachi Data Systems website.

## Microsoft SharePoint Server 2010

Microsoft SharePoint 2010 is a highly scalable platform for content management. For more information about Microsoft SharePoint Server 2010, see the [Microsoft SharePoint Server](#) web site.

## Microsoft SQL Server 2008 R2

The Microsoft website gives this overview for [Microsoft SQL Server 2008 R2](#):

Microsoft SQL Server 2008 R2 is a complete set of enterprise-ready technologies and tools that help people derive the most value from information at the lowest total-cost-of-ownership. Enjoy high levels of performance, availability, and security; employ more productive management and development tools; and deliver pervasive insight with self-service business intelligence (BI).

Microsoft SQL Server 2008 R2 also provides enhanced enterprise-class manageability for large OLTP deployments, like the one that is described in this reference architecture guide. When used with the Hitachi Virtual Storage Platform, Microsoft SQL Server 2008 R2 provides a scalable, high-performance database engine for any midrange to enterprise level application.

## Brocade Storage Area Network Switches

Two Brocade 5340 switches were used in the environment. These switches provide 8Gb/sec connectivity between the physical servers and the Hitachi Virtual Storage Platform, and 48 to 80 ports per switch. For more information, see the [Brocade and Hitachi Data Systems Products and Solutions web site](#).

## Solution Design

This is detailed information on the large scale Microsoft SharePoint 2010 environment design used for this reference architecture document. Although it includes application and hardware design information, it focuses primarily on the storage design required to build the infrastructure for the SharePoint environment.

## High-level Infrastructure

This solution environment was built with two base parameters in mind:

- Support for 200,000 users in a collaboration environment with a one percent user concurrency rate
- Availability of 20 individual site collections hosted under a single site

The architecture is built with the following:

- Microsoft Windows 2008 R2 Hyper-V as the hypervisor for all of the redundant SharePoint farm server infrastructure
- Hitachi Dynamic Provisioning and Hitachi Dynamic Tiering as the storage deployment technology of choice

These choices enabled this environment to be deployed quickly and managed effectively while maintaining the levels of high availability and performance required for Microsoft SharePoint deployments.

Figure 5 shows the Microsoft SharePoint infrastructure SQL and Hyper-V storage, and the Fiber Channel connectivity diagram. The term “PTD” in this diagram refers to the use of pass-through disks in Hyper-V for the guest virtual machines.

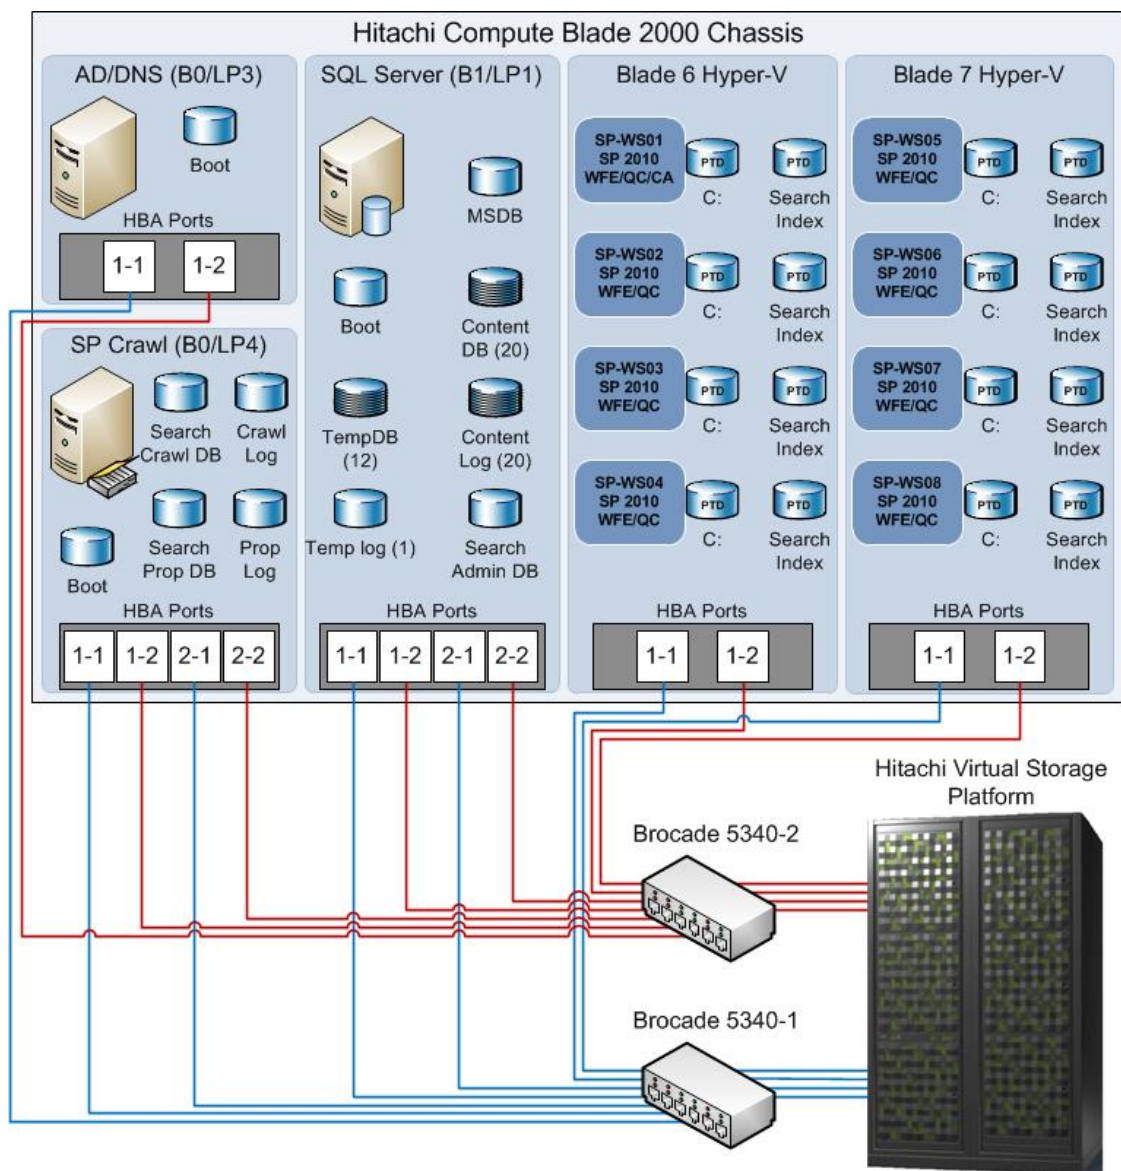


Figure 2

## Storage Architecture

This section describes the storage architecture created for the Microsoft SharePoint environment deployed for this reference architecture. It takes into consideration Hitachi Data Systems and Microsoft best practices for the deployment of large SharePoint environments.

### Storage Pools

This reference architecture uses three dynamic provisioning storage pools on a Hitachi Virtual Storage Platform. The three pools were:

- **HDT-MSFT-SP-DB(MT) (30)**—For the SQL database files (.mdf and .ndf), a dynamic provisioning pool using Hitachi Dynamic Tiering was created. It had two-tiers:
  - **Tier 1**—Used SAS drives to house frequently accessed data
  - **Tier 2**—Used cost-effective SATA drives to house the remaining data
- **HDT-MSFT-SP-LOG (31)**—This dynamic provisioning pool was used for the SQL database transaction log files (.ldf). It used SAS drives
- **HDT-MSFT-SP-VMPOOL (35)**—This dynamic provisioning pool was used for the guest virtual machines boot and index volumes. It used SAS drives.

Table 3 provides details on the specific pool configurations.

**Table 1. Dynamic Provisioning Pools**

<i>Dynamic Provisioning Pool Name (ID)</i>	<i>RAID Group Configuration</i>	<i>Drive Type</i>	<i>Number of RAID Groups</i>	<i>Pool Capacity (TB)</i>
HDT-MSFT-SP-DB(MT) (30)	RAID-5 (7D+1P)	300GB 10K RPM SAS	1	1.87
		2TB 7.2K RPM SATA	2	2.57
HDT-MSFT-SP-LOG (31)		300GB 10K RPM SAS	1	1.87
HDT-MSFT-SP-VMDATA (35)		300GB 10K RPM SAS	1	1.87

Table 4 through Table 11 provide details about the LDEVs used by the various hosts.

All volumes are thin provisioned on the Hitachi Virtual Storage Platform. The volumes were quick formatted in Microsoft Windows to avoid allocating all of the storage capacity configured for the volume. This tier arrangement accommodates any performance requirement for the SQL databases in our farm, with capacity for future growth.

Table 4 lists the volumes provisioned for the configuration database.

**Table 2. Configuration Database Volumes**

<i>Dynamic Provisioning Pool Name (ID)</i>	<i>LDEV</i>	<i>Size (GB)</i>	<i>Purpose</i>	<i>Storage Port</i>
HDT-MSFT-SP-DB(MT) (30)	30:00	10	Database	3A/4A/7A/8A
HDT-MSFT-SP-LOG (31)	31:00	1	Log	3A/4A/7A/8A

Table 5 lists the volumes provisioned for the central administration database.

**Table 3. Central Administration Database Volumes**

<i>Dynamic Provisioning Pool Name (ID)</i>	<i>LDEV</i>	<i>Size (GB)</i>	<i>Purpose</i>	<i>Storage Port</i>
HDT-MSFT-SP-DB(MT) (30)	30:01	10	Database	3A/4A/7A/8A
HDT-MSFT-SP-LOG (31)	31:01	1	Log	3A/4A/7A/8A

Table 6 lists the volumes provisioned for the content databases.

**Table 4. Content Databases Volumes**

<i>Dynamic Provisioning Pool Name (ID)</i>	<i>LDEVs</i>	<i>Size (GB)</i>	<i>Purpose</i>	<i>Storage Port</i>
HDT-MSFT-SP-DB(MT) (30)	30:02, 30:03, 30:04, 30:05, 30:06, 30:07, 30:08, 30:09, 30:0A, 30:0B, 30:0C, 30:0D, 30:0E, 30:0F, 30:10, 30:11, 30:12, 30:13, 30:14, 30:15	200	Databases 00-19	3A/4A/7A/8A
HDT-MSFT-SP-LOG (31)	31:02, 31:03, 31:04, 31:05, 31:06, 31:07, 31:08, 31:09, 31:0A, 31:0B, 31:0C, 31:0D, 31:0E, 31:0F, 31:10, 31:11, 31:12, 31:13, 31:14, 31:15	40	Logs 00-19	3A/4A/7A/8A

Table 7 lists the volumes provisioned for the search administration database.

**Table 5. Search Administration Volumes**

<i>Dynamic Provisioning Pool Name (ID)</i>	<i>LDEV</i>	<i>Size (GB)</i>	<i>Purpose</i>	<i>Storage Port</i>
HDT-MSFT-SP-DB(MT) (30)	30:16	10	Database	3A/4A/7A/8A
HDT-MSFT-SP-LOG (31)	31:16	2	Log	3A/4A/7A/8A

Table 8 lists the volumes provisioned for the crawl database

**Table 6. Crawl Database Volumes**

<i>Dynamic Provisioning Pool Name (ID)</i>	<i>LDEV</i>	<i>Size (GB)</i>	<i>Purpose</i>	<i>Storage Port</i>
HDT-MSFT-SP-DB(MT) (30)	30:17	185	Database	3A/4A/7A/8A
HDT-MSFT-SP-LOG (31)	31:17	36	Log	3A/4A/7A/8A

Table 9 lists the volumes provisioned for the search property database out of each dynamic provisioning pool.

**Table 7. Search Property Database Volumes**

<i>Dynamic Provisioning Pool Name (ID)</i>	<i>LDEV</i>	<i>Size (GB)</i>	<i>Purpose</i>	<i>Storage Port</i>
HDT-MSFT-SP-DB(MT) (30)	30:18	60	Database	3A/4A/7A/8A
HDT-MSFT-SP-LOG (31)	31:18	15	Log	3A/4A/7A/8A

Table 10 lists the volumes provisioned for the `tempdb` file out of each dynamic provisioning pool.

**Table 8. tempDB Volumes**

<i>Dynamic Provisioning Pool Name (ID)</i>	<i>LDEV</i>	<i>Size (GB)</i>	<i>Purpose</i>	<i>Storage Port</i>
HDT-MSFT-SP-DB(MT) (30)	30:50, 30:51, 30:52, 30:53, 30:54, 30:55, 30:56, 30:57, 30:58, 30:59, 30:5A, 30:5B	120	Databases 0-11	3A/4A/7A/8A
HDT-MSFT-SP-LOG (31)	30:31	320	Log	3A/4A/7A/8A

Table 11 lists the volumes provisioned for the Hyper-V VMs and Index volumes.

**Table 9. Virtual Machine Volumes**

<i>Dynamic Provisioning Pool Name (ID)</i>	<i>LDEV</i>	<i>Size (GB)</i>	<i>Purpose</i>	<i>Storage Port</i>
HDT-MSFT-SP-VMPOOL (35)	30:20, 30:24	150	SP-WS01 boot Vol. SP-WS01 IDX vol	3B/7B
HDT-MSFT-SP-VMPOOL (35)	30:21, 30:25	150	SP-WS02 boot Vol. SP-WS02 IDX vol	3B/7B
HDT-MSFT-SP-VMPOOL (35)	30:22, 30:26	150	SP-WS03 boot Vol. SP-WS03 IDX vol	3B/7B
HDT-MSFT-SP-VMPOOL (35)	30:23, 30:27	150	SP-WS04 boot Vol. SP-WS04 IDX vol	3B/7B
HDT-MSFT-SP-VMPOOL (35)	30:28, 30:2C	150	SP-WS05 boot Vol. SP-WS05 IDX vol	3B/7B
HDT-MSFT-SP-VMPOOL (35)	30:29, 30:2D	150	SP-WS06 boot Vol. SP-WS06 IDX vol	3B/7B
HDT-MSFT-SP-VMPOOL (35)	30:2A, 30:2E	150	SP-WS07 boot Vol. SP-WS07 IDX vol	3B/7B
HDT-MSFT-SP-VMPOOL (35)	30:2B, 30:2F	150	SP-WS08 boot Vol. SP-WS08 IDX vol	3B/7B

## Capacity Planning

The capacity and performance requirements for a Microsoft SharePoint environment are dictated by many variables, such as:

- The number of documents and their size
- The type of content being saved
- The number of users that will use the environment concurrently

The capacity required by your architecture is dictated primarily by the amount of content that will be stored within Microsoft SharePoint and the web site structure of your environment. With Hitachi Dynamic Provisioning and Hitachi Dynamic Tiering, all RAID groups that are wide striped in each tier must be of the same RAID type and drive type.

Keep the following Microsoft SharePoint sizing recommendations and limits in mind:

- Limit each content database to 200GB. As an exception, limit each content database used for a static archive to 1TB.
- If you plan to use Microsoft SharePoint Server 2010 granular backup tools, limit each site collection to 100GB.
- The minimum size of the crawl database is 4.6 percent of the sum of the sizes of each content database.
- The minimum size of the search property database is 1.5 percent of the sum of the sizes of each content database.
- The minimum size of the configuration database is 2GB. It grows by approximately 40MB for each 50,000 site collections.
- The minimum size of the central administration content database is 1GB.
- The sum of the index volumes is at least 3.5 percent of the sum of the sizes of each content databases multiplied by four (to account for index repartitioning).
- In most cases, the minimum size of transaction logs for a database is 20 percent the size of the database.
- Use the simple recovery model for the configuration database because it tends to become large over time.

For Microsoft SharePoint 2010 content databases that exceed 100GB, provide a fixed number in megabytes for the auto-growth value instead of a percentage on the SQL server. This reduces how often these file sizes must be increased, which minimizes the possible performance overhead on the server and the storage system caused by the increase in file sizes. Set **MAXSIZE** on each database file to a value that matches the capacity of the volume being used. You may add volumes and files to the database later, if needed.

Consider the additional storage requirements if you plan to install additional features in your Microsoft SharePoint environment. This reference architecture guide only focuses on the key components that an enterprise-class SharePoint environment must have to support a collaboration environment.

For more information about storage capacity planning for a Microsoft SharePoint deployment, see the Microsoft *TechNet* article "[Storage and SQL Server capacity planning and configuration \(SharePoint Server 2010\).](#)"

## Architecture IOPS Planning

When determining your IOPS requirement, keep these Microsoft guidelines in mind:

- The crawl database requires between 3,500 and 7,000 IOPS.
- The search property database requires about 2,000 IOPS.

Also keep these Microsoft SQL Server recommendations in mind:

- Create dedicated LUNs for each of the databases used for the Microsoft SharePoint environment. This provision defines a unique file system and IO queue for each SQL data file. Due to the number of data files, databases and log LUNs were mapped to mount points using drive C as the anchor point.
- Pre-size the data and log files rather than relying on auto-growth. You may enable auto-growth to meet deployment needs for safety purposes as long as you manage the growth of the data files with proper planning. In addition, turn on **MAXSIZE** for each database file to ensure none of them grow to a point where it uses the entire space available on a disk.
- Create one `tempDB` data file per CPU core and make each `tempDB` file equal in size. 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 Microsoft *Developer Network* article "[Optimizing tempDB Performance](#)."

No precise method exists to establish content database IOPS requirements. Like any other SQL server environment, the recommended practice is first to deploy a test environment to establish whether your content databases can support your IOPS requirement.

## Host Considerations

To ensure optimal performance from your storage system, understand the following during the planning process:

- NTFS allocation unit size
- Path redundancy
- Host bus adapter (HBA) placement
- Queue depth or execution throttle setting in the HBA driver

### *NTFS Allocation Unit Size*

Choose an allocation unit size that matches the I/O profile of your application. This allows more efficient I/O on your storage system. Hitachi Data Systems recommends, when formatting your SQL volumes in Microsoft Windows Server, overriding the default setting and specify an allocation unit size of 64KB.

### *Host Bus Adapter Drivers*

When choosing host bus adapter (HBA) drivers, ensure that you use the most current recommended drivers for the Hitachi Virtual Storage Platform. Major HBA vendors allow you to download current drivers for Hitachi storage systems. For a list of currently supported HBAs and drivers, see [Interoperability](#) on the Hitachi Data Systems web site.

### *General Database Maintenance*

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

- Monitor the database server to make sure that it responds 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. It 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.

For more information, see the Microsoft *TechNet* article "[Monitoring and maintaining SharePoint Server 2010.](#)"

## SharePoint Architecture

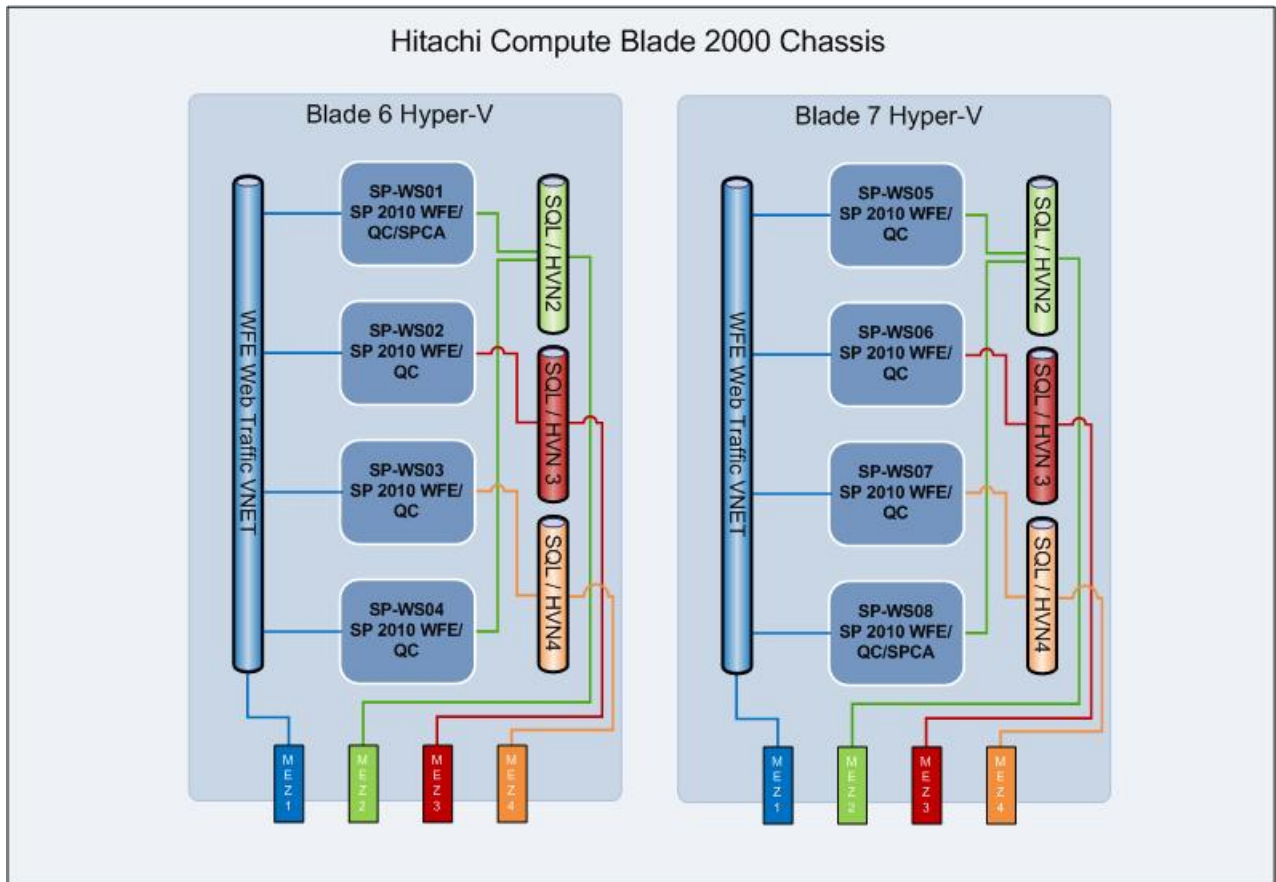
This solution provides the basic functionalities for a large collaboration environment by using the following:

- Eight web servers with the network load balance feature enabled. Each web server also contained a separate query component, a local search index component, and a mirror of one of the other web server index components.
- One SQL server for the Microsoft SharePoint farm.
- One search crawl server with a dedicated SQL instance that provides the crawling functionality to the farm.

Microsoft SharePoint farms scale out with the addition of servers to the environment in order to accommodate additional users and additional site content. Servers can be physical or virtual. Servers can be added to or removed from the farm, as needed.

A mirrored database server was not specified in this guide. As a best practice, Microsoft strongly recommends the use of database mirroring for the SharePoint database servers to provide high availability for the critical data back end, the foundation of any SharePoint environment. For more information about setup and requirements for SQL server database mirroring, see the [Database Mirroring Setup Overview](#) on Microsoft *TechNet*.

Figure 6 shows the networking configuration for the internal networking of the Microsoft Hyper-V hosts.



**Figure 3**

There are two different types of switches available for the Hitachi Compute Blade chassis. They are differentiated by the external ports arrangement. The switch type installed in this chassis for the mezzanine four port gigabit cards includes two 10G ports. One switch was utilized as a trunk between those switches. The mezzanine ports were utilized exclusively for this reason.

The number of web servers that your environment requires depends primarily on the amount of load the users put on the environment due to the number of HTTP requests. Carefully monitor each web server's resources to ensure that they are capable of handling users' requests. The primary resources to monitor in comparison to their limits are the following:

- CPU
- Memory
- Network utilization rates

For example, if the CPU utilization rate averages at about an 80 percent level during regular business hours for all of the web servers currently in the environment, then consider adding more web servers to the environment.

## SQL Server Architecture

When deploying Microsoft SQL Server in support of the SharePoint databases, understand the essential key file types for a database. Understanding each file type's workload, along with the database type and size, enables storage and database administrators (DBAs) to establish storage requirements for the SQL server and SharePoint environment.

### SQL Server Database files

These databases were defined in a single Hitachi Dynamic Provisioning pool:

- SQL Server administrative
- SharePoint administrative
- SharePoint content
- SharePoint search

Each database was defined with an individual LUN, with the exception of the SQL server administrative databases.

### SQL Server Transaction Log File

Every SQL Server database has at least one log volume and file that records database modifications made by each transaction. The log volume and files are critical components of the database for availability and data integrity.

In the event of a system failure, at a minimum the active transaction log is required to bring the database back to a consistent state. It writes the transaction logs before updating the data records to the database file via the checkpoint process. Use the logs to roll back transactions if you discover corruption or to recover a committed transaction if an error writing to the database occurs.

Microsoft suggests aiming for log I/O response times between 1 and 5 milliseconds. Microsoft and Hitachi Data Systems recommend provisioning separate dedicated volumes for database and transaction log files. If you design your storage properly, it writes the transaction logs before updating the data records to the database file via the checkpoint process.

One of the features built into the Hitachi Virtual Storage Platform is the ability to optimize physical I/O based on recognition of I/O patterns. The controller of the Virtual Storage Platform storage system 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 Virtual Storage Platform can significantly increase overall performance.

### SQL Server *tempDB* Files

SQL Server *tempDB* files are used for storage of temporary data structures. They 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 at Hitachi Data Systems shows that it can reside on the same pool as the databases without adverse effects. Accordingly, if the environment does not have sufficient physical disk I/O resources to meet the combined requirements of `tempDB` and 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.

By default, `tempDB` only supports a single data file group and a single log file group with a default number of files set to 1. Microsoft recommends creating at least as many data files of equal size as you have CPU cores. The number of concurrent threads is less than or equal to the number of CPU cores.

If the `tempDB` file does not perform well, the system does not perform well. New allocations favor the newer file because of availability free space. This can cause allocation problems. Restart SQL server when you add a new file.

### *Database Performance*

Performance considerations include monitoring and measuring average disk read and write response times. The average response time, measured in milliseconds for a logical or physical disk, can be affected by I/O size, RAID configuration, and other factors in the data path.

Disk latency response times can vary, depending on the size and nature of I/Os being issued. These numbers vary across different storage configurations. Cache size and utilization can greatly affect this.

In most cases, 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 file group architectures to improve performance, operational convenience, or availability. Acceptable database I/O response times typically range between 5 and 20 milliseconds.

In this solution, scaling the search database as content growth increased was addressed by the following:

- Adding additional LDEVs to the pool when they became necessary from a capacity or performance perspective
- Creating a second data file of equal size for expansion

For more information about SQL files and file groups, see the Microsoft Enterprise Search Blog post "[SQL File group and Search.](#)"

## Server Architecture

The server being used for the Microsoft SharePoint Server infrastructure plays a key role in the performance and availability levels for your environment. When choosing servers for your environment, take into consideration not only Microsoft's minimum and recommend values, but also the overall I/O bandwidth architecture of the server being used.

Two Hitachi Compute Blade 2000 server blades were used in this solution as Microsoft Hyper-V hosts for the virtual machines in the SharePoint environment. To better serve the performance needs of the environment, two additional server blades were divided up using logical partitions for the SQL server, search crawl, and Active Directory.

Through the use of logical partitions, you can asymmetrically divide server blades for multiple roles requiring different amounts of resources. In this reference architecture, the server blades had larger partitions for the SQL server and search crawl servers, and a smaller partition for Active Directory.

The SQL server required 12 processor cores, and 96GB of memory. Neither of these requirements could have been satisfied by a Hyper-V VM, so a logical partition was used. The remaining processor and memory resources in each blade were used for other logical partitions serving other roles, allowing full usage of each blade.

---

**Note**— Even though the Hitachi Compute Blade 2000 contains industry standard PCIe slots, when using logical partitions, the use of Hitachi HBAs is required.

---

Table 12 shows the server configuration.

**Table 10. Server Configuration**

<i>Server</i>	<i>Description</i>	<i>Firmware Level</i>	<i>Qty</i>	<i>Operating System</i>
Microsoft Hyper-V host servers	Hitachi Compute Blade 2000 E55A2 <ul style="list-style-type: none"> <li>• Two Intel Xeon X5670 at 2.93GHz</li> <li>• 72GB RAM</li> <li>• Two Emulex LPe 12002-M8 HBA</li> </ul>	EFI BIOS v4.6.3	2	Microsoft Windows Server 2008 R2 Datacenter Edition
Microsoft SQL Server	Hitachi Compute Blade 2000 E57A1 <ul style="list-style-type: none"> <li>• Two Intel Xeon X7560 at 2.26GHz</li> <li>• Logical partition</li> <li>• 12 Cores, 96GB RAM</li> <li>• Two Hitachi 8Gb/sec dual port HBA</li> </ul>	EFI BIOS v4.6.3	1	Windows Server 2008 R2 Enterprise Edition
Microsoft SharePoint Search Crawl— sharing server blade with Active Directory and DNS	Hitachi Compute Blade 2000 E57A1 <ul style="list-style-type: none"> <li>• Two Intel Xeon X7560 at 2.26GHz</li> <li>• Logical partition</li> <li>• 4 Cores, 32GB RAM</li> <li>• One Hitachi 8Gb/sec dual port HBA</li> </ul>	EFI BIOS v4.6.3	1	Windows Server 2008 R2 Enterprise Edition

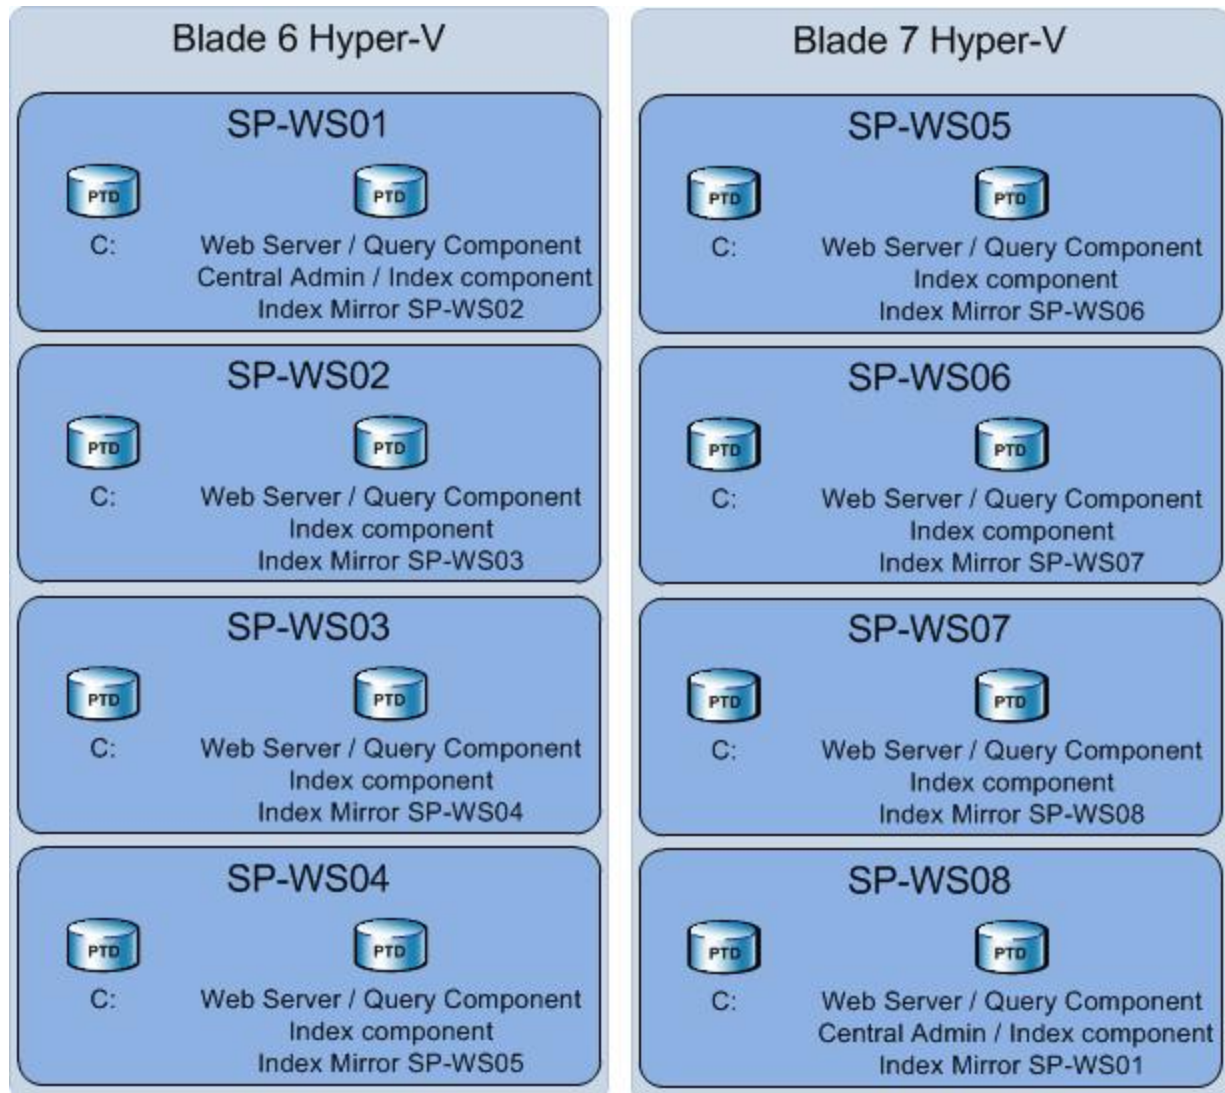
<i>Server</i>	<i>Description</i>	<i>Firmware Level</i>	<i>Qty</i>	<i>Operating System</i>
Active Directory and DNS—sharing Blade with Microsoft SharePoint Search Crawl	Hitachi Compute Blade 2000 E57A1 <ul style="list-style-type: none"> <li>• Two Intel Xeon X7560 at 2.26GHz</li> <li>• Logical partition</li> <li>• 4 Cores, 32GB RAM</li> <li>• One Hitachi 8Gb/sec dual port HBA</li> </ul>	EFI BIOS v4.6.3	1	Windows Server 2008 R2 Enterprise Edition

Table 13 lists the virtual servers used in this reference architecture.

**Table 11. Virtual Server Configuration**

<i>Name</i>	<i>Role</i>	<i>Server</i>	<i>Qty</i>	<i>Drives</i>	<i>Details</i>
SP-WS01 SP-WS02 SP-WS03 SP-WS04	Web Front Ends, Search Index, and Mirror	Blade6-HV	4	150GB LUN pass-through disk for boot for each VM 150GB LUN pass-through disk for index data for each VM	Windows 2008 R2 Enterprise Edition, 4 VCPUs, 8GB RAM
SP-WS05 SP-WS06 SP-WS07 SP-WS08	Web Front Ends, Search Index, and Mirror	Blade7-HV	4	150GB LUN pass-through disk for boot for each VM 150GB LUN pass-through disk for index data for each VM	Windows 2008 R2 Enterprise Edition, 4 VCPUs, 8GB RAM

Figure 4 details the Microsoft SharePoint role services, index partitions, and mirrors on the Microsoft Hyper-V hosts.



**Figure 4**

Microsoft Hyper-V virtual machines can access LUNs in two ways:

- Pass-through disks
- Virtual hard disks (VHDs) created in a volume.

Pass-through disks enable the LUNs to be available to a virtual machine as if they were directly connected to it. All of the volumes used by the guest virtual machines in this reference architecture are pass-through disks, presented from the storage system.

For more information about virtualization planning for SharePoint environments, see the Microsoft TechNet article "[Virtualization Planning \(SharePoint Server 2010\).](#)"

## SAN Architecture

The LDEVs provisioned for the Microsoft SharePoint environment were mapped to multiple ports on the Hitachi Virtual Storage Platform. All switch zoning defined with single initiator to target ports on the Hitachi Virtual Storage Platform.

LDEVs were assigned to the ports in a way that ensures that the host had two or more paths to the storage system for high availability and load balancing performance.

- The SQL server used four HBA ports, enabling a four-path connection for all LUNs mapped to that server. This allows The SQL server to load balance the connections. The SQL server used four ports on the Virtual Storage Platform (3A, 4A, 7A, 8A).
- The Microsoft Hyper-V virtual machine guest drive C and index volumes used two ports on the Virtual Storage Platform (3B, 7B).
- The SharePoint 2010 search crawler SQL server used four ports on the Virtual Storage Platform (3A, 4A, 7A, 8A).
- The logical partitions (SQL, SP-WS-SC, and DC3) have their OS boot volumes on the Virtual Storage Platform which used one additional port on the VSP (4B).

The environment used two Brocade 5340 switches to provide scalability and high availability to the environment.

Table 14 lists the zoning details for the SAN.

**Table 12. SAN Switch Architecture**

<i>Server</i>	<i>HBA Ports</i>	<i>Switch Zone</i>	<i>Storage Port</i>	<i>Switch</i>
Blade6-HV	HBA1-1	BS2K_01_B6_HBA1_1_ASE45_36_3B	3B	5300-01
Blade6-HV	HBA2-2	BS2K_01_B6_HBA2_2_ASE45_36_7B	7B	5300-02
Blade7-HV	HBA1-1	BS2K_01_B7_HBA1_1_ASE45_36_3B	3B	5300-01
Blade7-HV	HBA2-2	BS2K_01_B7_HBA2_2_ASE45_36_7B	7B	5300-02
MSFT-SP-SQL-1	HBA1-1	BS2K_01_B1_L1_HBA1_1_ASE45_36_3A	3A	5300-01
MSFT-SP-SQL-1	HBA1-2	BS2K_01_B1_L1_HBA2_1_ASE45_36_4A	4A	5300-01
MSFT-SP-SQL-1	HBA2-1	BS2K_01_B1_L1_HBA1_2_ASE45_36_7A	7A	5300-02
MSFT-SP-SQL-1	HBA2-2	BS2K_01_B1_L1_HBA2_2_ASE45_36_8A	8A	5300-02
MSFT-SP-WS-SC	HBA1-1	BS2K_01_B0_L4_HBA1_1_ASE45_36_3A	3A	5300-01
MSFT-SP-WS-SC	HBA1-2	BS2K_01_B0_L4_HBA1_2_ASE45_36_4A	4A	5300-01
MSFT-SP-WS-SC	HBA2-1	BS2K_01_B0_L4_HBA2_1_ASE45_36_7A	7A	5300-02
MSFT-SP-WS-SC	HBA 2-2	BS2K_01_B0_L4_HBA2_2_ASE45_36_8A	8A	5300-02
Domain Controller	HBA1-1	BS2K_01_B0_L3_HBA1_1_ASE45_36_4B	4B	5300-01
SANBOOT MSFT-SP-SQL-1	HBA1-1	BS2K_01_B1_L1_HBA1_1_ASE45_36_4B	4B	5300-01
SANBOOT MSFT-SP-WS-SC	HBA1-1	BS2K_01_B0_L4_HBA1_1_ASE45_36_4B	4B	5300-01

When designing your SAN architecture, follow these best practices to ensure a secure, high-performance and scalable Microsoft SharePoint deployment:

- Use at least two HBAs, and place them on different buses within the server. This distributes the workload over the server's PCIe bus architecture. Dual-port HBAs accomplish this by dividing the bandwidth of each port between PCIe lanes on the expansion bus.
- Use dual SAN fabrics, multiple HBAs, and host-based multi-pathing software when using Microsoft SQL Server in a business-critical deployment. Two or more paths from the SQL and application servers connecting to two independent SAN fabrics are essential to ensure the redundancy required for critical applications.
- Zone your fabric appropriately for multiple, unique paths from HBAs to storage ports. Use single initiator zoning. Use at least two Fibre Channel switch fabrics to provide multiple independent paths to the Hitachi Virtual Storage Platform to prevent configuration errors from disrupting the entire SAN infrastructure.
- For large bandwidth requirements that surpass a single HBA's port capability, use additional HBAs and use the round robin load-balancing setting for Hitachi Dynamic Link Manager or Microsoft Windows Server 2008 R2 native MPIO.

## Physical and Virtual Network Architecture

Following best practice from Microsoft, the Microsoft Hyper-V hosts in this reference architecture have separate 1Gb/sec physical network adapters to be used independently for different types of data traffic.

- One adapter is used primarily for Hyper-V host data traffic for management purposes.
- The other four adapters are dedicated for virtual machine traffic.

This configuration provides improved performance and adapter redundancy. Each Hyper-V host is configured with four virtual LANs that provide the network connections to the virtual machines.

- NLB web traffic was restricted to one dedicated virtual network.
- SQL and management traffic was load balanced across the remaining three networks.

Each virtual machine was configured with two network adapters (see Figure 6).

- One adapter provides a dedicated link between the servers as well as for management purposes.
- Another adapter provides a dedicated channel between the users and the environment servers (primarily the web servers).

Network load balancing was installed on all web servers and configured to use the network interface dedicated for user connections to the environment.

Figure 8 shows the network configuration for the test environment

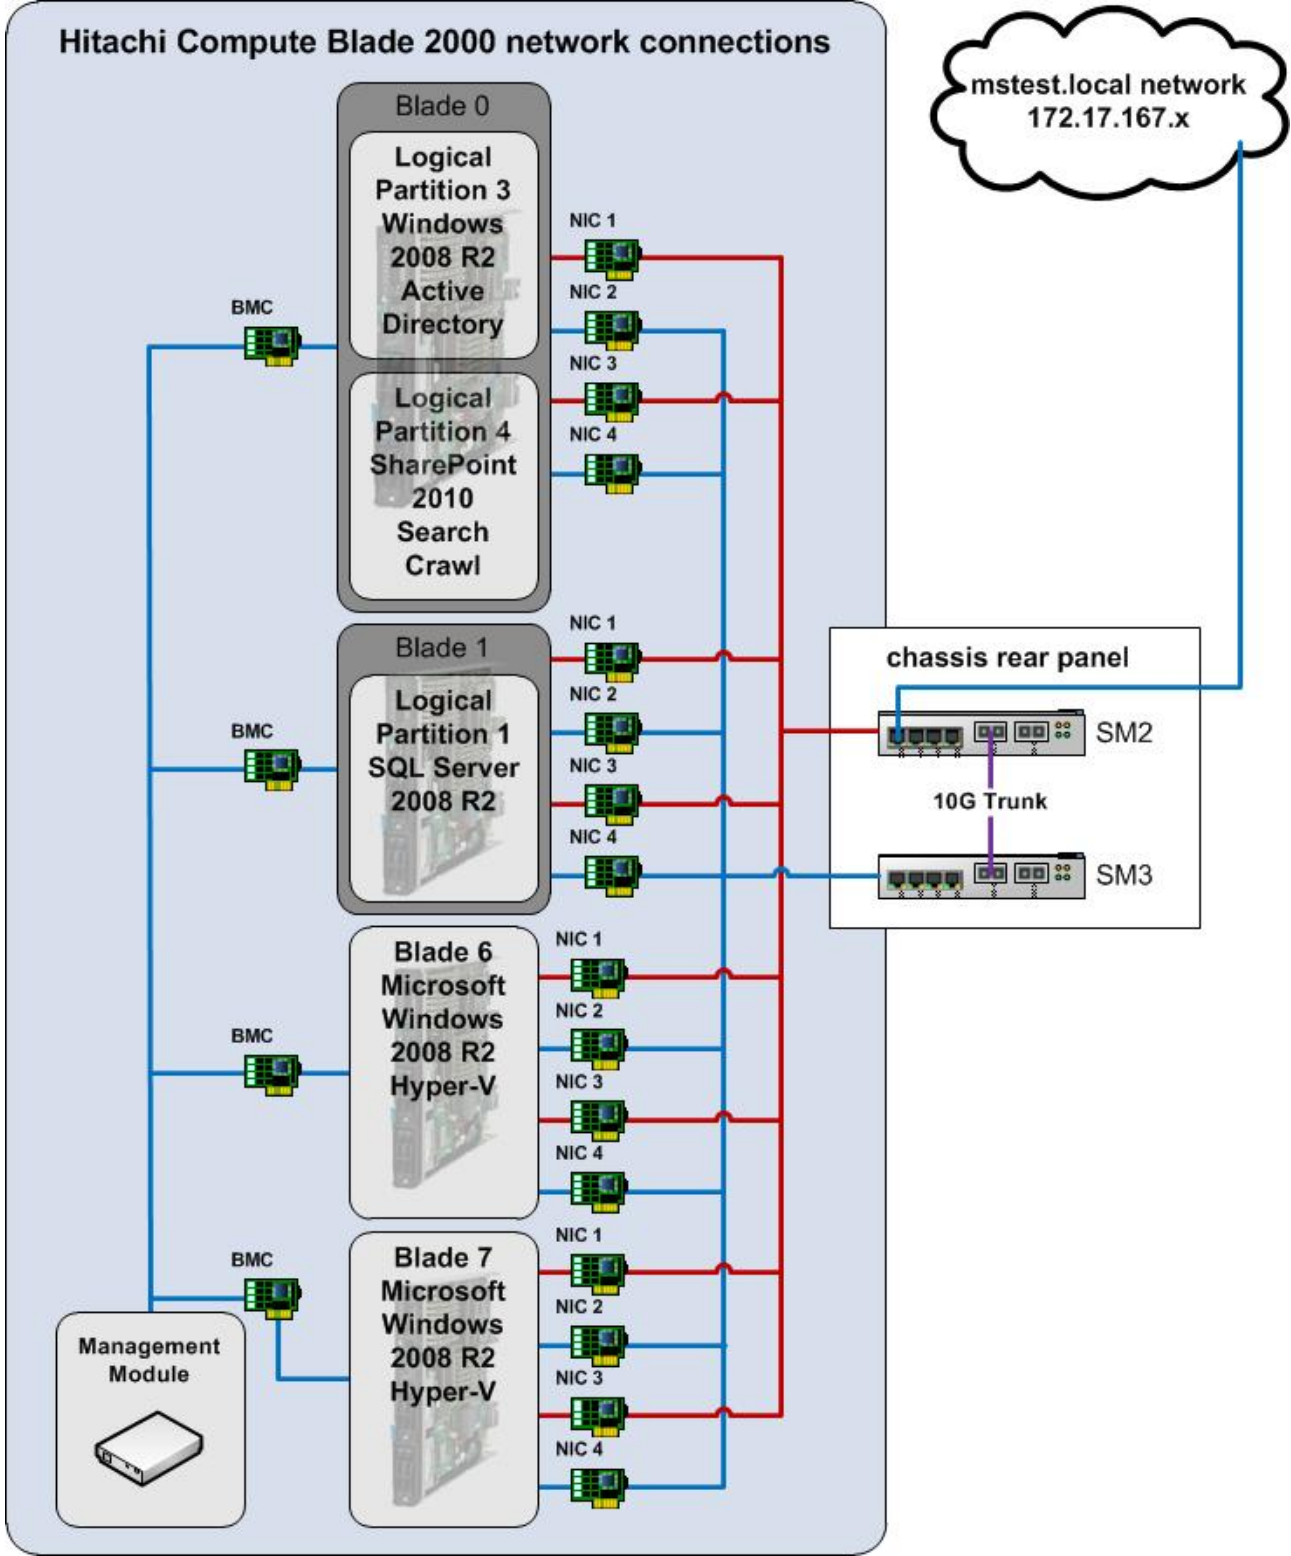


Figure 8

# Engineering Validation

This describes the testing methodology and results that validate the reference architecture described in this white paper.

## Test Methodology

Hitachi Data Systems used Microsoft Visual Studio 2010 Ultimate edition with its Visual Studio test agents to test this reference architecture. Different load tests created a user load that is similar to what most companies would see on a production Microsoft SharePoint 2010 environment. Table 15 describes the user load mix that was used during testing.

**Table 13. User Load Type Mix**

<i>Load Type</i>	<i>Percentage</i>
Browsing	75
Upload	3
Search	11
Open	11

For these tests, the Visual Studio host, the test agent, and test controller hosts were in a separate network dedicated for user data connectivity to the web servers in the SharePoint environment. Six test agents were used to evenly drive the load mix to all 20 site collections in the environment.

A total of 200,000 users were created and given contributor permission to the environment's site collections. A one percent user concurrency rate was set, meaning that at any given time, 2,000 users were connected to the SharePoint environment using the load mix described in Table 15.

Data population and growth utilized real world files and documents of different types and sizes to ensure the closest similarity to real world environments. The files used to populate the content databases and test the environment had variable size ranging from a couple kilobytes to a couple megabytes. The 100MB default limitation was maintained for upload size.

The content databases were attached to each of 20 site collections, and were loaded with about 60GB of these files. There was also a repository set aside of these random files for the upload tests. During the collaboration test execution, the data files were randomly uploaded to all of the site collections by random test users.

The test methodology's primary function was to ensure that the environment could withstand a 1 percent user concurrency and execution of a mix load type while maintaining acceptable levels of response time based on the calculation of a heavy user profile. This means that overall the test should achieve and maintain an average of 33 requests per second. Table 16 lists acceptable page response times that were used for the testing of the environment.

**Table 14. Acceptable Response Times**

<i>Load Type</i>	<i>Acceptable Page Response Time</i>
Browsing	< 3sec
Upload	Variable (dependent on file size)
Search	< 5sec (using default SharePoint search page)
Open	Variable (dependent on file size)

For the search tests, the SharePoint search page was utilized to execute searches against the different content databases as well as the overall SharePoint environment. Because documents from 10KB to 5MB were used, the acceptable response times for certain load types were variable. The average response time for the entire load remained within the acceptable range expected for the difference file sizes. A modify workload type can be accounted for by using the values for open and upload, as the overall act of modifying a file by individual users will depend on the amount of changes done by the user.

## Test Results

Although the primary function of the test was to ensure that the response times for each load type remained at or below acceptable values, the environment resources were monitored during testing to ensure that there were not bottlenecks. The test results show that the environment successfully handled the user load within the acceptable response time values.

Table 17 lists the response times for each of the load types.

**Table 15. Response Times Results**

<i>Load Type</i>	<i>Response Time</i>
Browsing	< 3sec
Upload	Variable $\leq$ 6sec
Search	< 4sec
Open	Variable $\leq$ 3sec

Table 18 lists the average and maximum CPU utilization rates for each of the virtual servers utilized in the environment. The CPU utilization rate for the mirrored SQL server was not monitored.

**Table 16. Average and Maximum CPU Utilization Rates**

<i>Type</i>	<i>Average CPU Utilization Rate</i>	<i>Maximum CPU Utilization Rate</i>
Production SQL Server Host	67.9%	98.9%
Web Servers (overall average)	35.0%	70.0%
Crawl Servers	6.0%	21.0%

## Conclusion

This document describes a tested reference architecture that shows the storage savings benefit of Hitachi Dynamic Tiering functionality when applied to a large Microsoft SharePoint 2010 farm.

When compared to a similar environment that did not use Hitachi Dynamic Tiering, there were the following benefits:

- Maintain the performance and service level agreements (SLA) when using a multi tier pool.
- A cost savings that resulted from the reduction of the total number of drives versus what would be needed with a conventional pool or dedicated RAID groups.
- A cost savings by enabling the use of less expensive SATA drives in the multi tier pool.

Hitachi Data Systems Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the [Hitachi Data Systems Global Services web site](#).

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the [Hitachi Data Systems Resources web site](#). Click the **Product Demos** tab for a list of available recorded demonstrations.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the [Hitachi Data Systems Education web site](#).

For more information about Hitachi products and services, contact your sales representative or channel partner or visit the [Hitachi Data Systems web site](#).

 **Hitachi Data Systems Corporation**

---

*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 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 Corporation*

© Hitachi Data Systems Corporation 2011. All Rights Reserved. AS-093-00 July 2011

**Corporate Headquarters**

750 Central Expressway,  
Santa Clara, California 95050-2627 USA  
[www.hds.com](http://www.hds.com)

**Regional Contact Information**

**Americas:** +1 408 970 1000 or [info@hds.com](mailto:info@hds.com)  
**Europe, Middle East and Africa:** +44 (0) 1753 618000 or [info.emea@hds.com](mailto:info.emea@hds.com)  
**Asia Pacific:** +852 3189 7900 or [hds.marketing.apac@hds.com](mailto:hds.marketing.apac@hds.com)