Deploying Microsoft SharePoint Server 2010 on the Hitachi Virtual Storage Platform

Reference Architecture Guide

By Eduardo Freitas

February 2011

Feedback

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

Table of Contents

Microsoft SharePoint Server 2010 on the Hitachi Virtual Storage Platform

Reference Architecture

Businesses increasingly rely on collaboration tools to enable employees to work together, share information and manage content across countries, continents and time zones. Microsoft SharePoint Server 2010 is often the tool of choice to achieve those important objectives. However, planning and deploying SharePoint in enterprise environments is a complex endeavor requiring many hours of effort. Creating a robust deployment requires complex calculations that take into account a number of factors. Hitachi simplifies that task with the Hitachi Virtual Storage Platform, an enterprise-class storage system, and Hitachi Dynamic Provisioning software, a thin provisioning and wide striping software product.

In today’s data-driven economy, information is the new currency. This information exists in many forms and must be stored in a manner that makes it readily accessible, and it must be protected to ensure an organization’s survival and success. This must all be done while maximizing cost efficiency and return on investment. The Hitachi Virtual Storage Platform can help companies achieve these goals by creating an agile storage infrastructure that reduces costs and increases performance, availability, scalability and reliability.

The Hitachi Virtual Storage Platform is the industry’s only 3D scaling storage platform. With the unique ability to concurrently scale up, scale out and scale deep in a single storage system, the Virtual Storage Platform flexibly adapts for performance, capacity, connectivity and virtualization. No other enterprise storage platform can dynamically scale in three dimensions. Scaling up allows you to increase virtual server consolidation, improve utilization of resources, and reduce costs. Scaling out allows you to meet increasing demands by combining multiple chassis into a single logical system with shared resources. Scaling deep extends the advanced functions of the Virtual Storage Platform to external multivendor storage.

This document describes a tested reference architecture that provides high availability and simplified storage administration for enterprise SharePoint deployments. It is written for storage administrators, SQL administrators and SharePoint administrators. Readers need to be familiar with general storage and SharePoint implementation concepts.

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. This white paper provides a tested, scale-up reference architecture that simplifies the planning process from a storage perspective. Hitachi Data Systems testing used two Windows Server 2008 R2 servers with the Hyper-V role and multiple virtual machines. The virtual machines were used to host the two search servers, two application servers, and ten web servers using network load balancing (NLB). Two physical servers were used to host the production SQL Server 2008 R2 instance and a mirrored SQL Server 2008 R2 instance. Each physical server was connected to the storage system through the use of two dual port HBAs to allow for multiple paths for both load balancing and failover capabilities. The test environment included a Windows Domain Controller supplying Active Directory and DNS services that was already part of the existing lab infrastructure.

The solution described in this paper was designed with a large organization in mind, initially supporting 200,000 users with 20 site collections (using 20 content databases) hosted on a single web application. The use of the latest technologies from both Microsoft and Hitachi Data Systems enables the effective use of resources available in your data center while maintaining high availability and performance levels required for a SharePoint environment. The SharePoint server infrastructure is virtualized using Hyper-V hypervisor, enabling better control over the server sprawl issue that is common in SharePoint environments. The storage infrastructure uses the Hitachi Virtual Storage Platform to provide the following:

- Robust, scalable architecture for storing, managing and accessing large and growing amounts of your organization's information hosted in SharePoint Server environments
- Consolidated files and storage to ease management burden and reduce operational overhead through the use of Hitachi Dynamic Provisioning software

Figure 1 provides a high-level view of the infrastructure used to host the SharePoint environment.

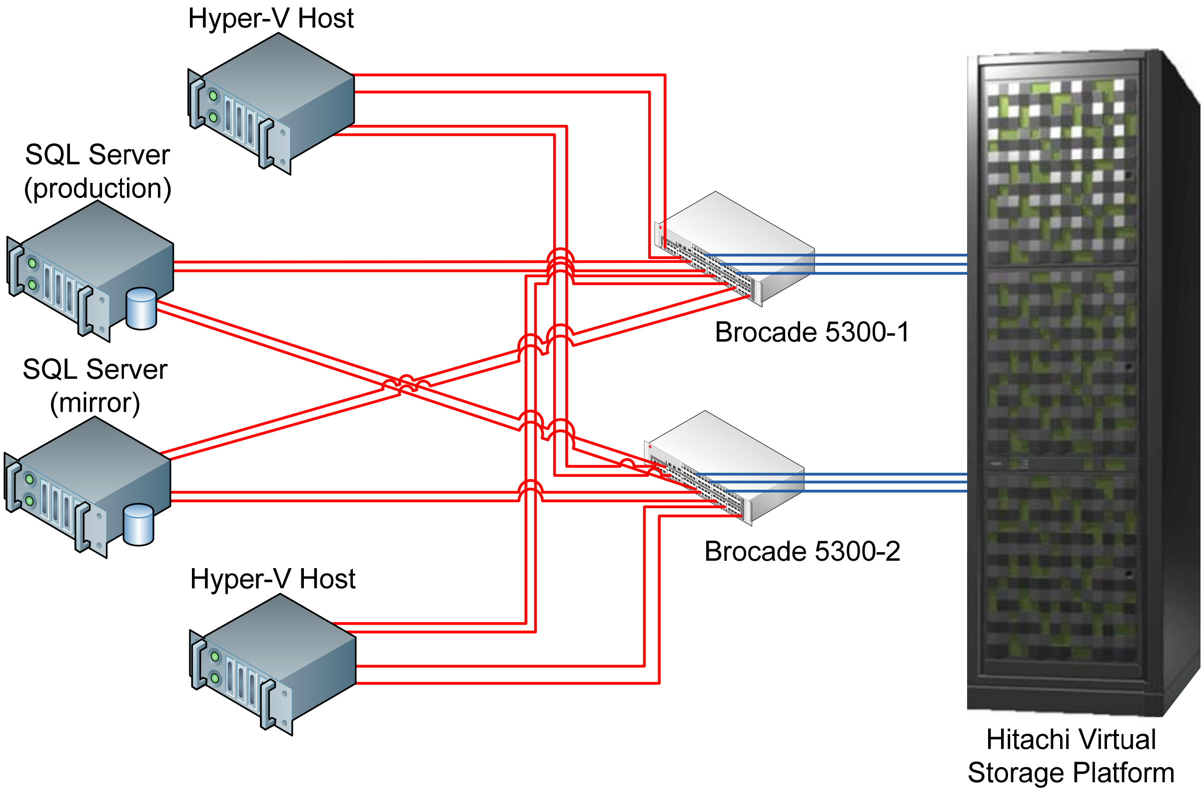
Hyper-V Host
SQL Server
(production)
SQL Server
(mirror)
Hyper-V Host
Brocade 5300-1
Brocade 5300-2
Hitachi Virtual
Storage Platform

Figure 1

This reference architecture enables storage and database administrators to successfully plan and deploy SharePoint environments by describing the interaction between the many components of SharePoint Server, SQL Server and how they utilize the Hitachi Virtual Storage Platform storage to create a robust and highly available SharePoint environment.

Key Solution Components

Whether on a single or multi server farm deployment, for a development, staging or production environment, SharePoint implementations all have the basic roles that are necessary for their functionality. This reference architecture addresses the deployment of a large scale production environment that uses dedicated servers for each of SharePoint's infrastructure components. Most importantly, the Hitachi Virtual Storage Platform provides storage for both SQL Server and index volumes for the environment.

The following sections describe each of the key components for the solution described in this paper.

Hitachi Virtual Storage Platform

The Hitachi Virtual Storage Platform is the industry's only 3D scaling storage platform. With the unique ability to concurrently scale up, scale out and scale deep in a single storage system, the new Virtual Storage Platform flexibly adapts for performance, capacity, connectivity and virtualization. No other enterprise storage platform can dynamically scale in three dimensions. The Virtual Storage Platform provides virtual storage that meets the growing demands of server virtualization.

The trend in server virtualization is to consolidate the I/O workload of many servers onto a single storage system. As more virtual machines are consolidated onto a physical host, storage systems must be able to dynamically add more storage resources to keep up with I/O demand. The 3D scaling capability of the Virtual Storage Platform meets that requirement.

Scaling up allows you to increase virtual server consolidation, improve utilization of resources, and reduce costs. With the Hitachi Virtual Storage Platform, you can increase performance, capacity and connectivity by adding cache, processors, connections and disks to the base system. A virtual server that accesses the storage system can use all these resources, which act as one system managed as a common pool of resources.

Scaling out allows you to meet increasing demands by combining multiple chassis into a single logical system with shared resources. By scaling out you can support increased resource needs in virtualized server environments.

Scaling deep extends the advanced functions of the Virtual Storage Platform to external multivendor storage. By dynamically virtualizing new and existing storage systems, those systems become part of the Virtual Storage Platform's pool of storage resources. Once virtualized, external data can then be migrated, tiered, replicated and managed by the Virtual Storage Platform. In this manner, older data storage systems can gain a longer useful life. You can extend distance replication for business continuity to lower-cost, lower-function storage systems by virtualizing them behind a Virtual Storage Platform.

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. You can add redundant pairs of directors and cache modules as required without disruption to connected host servers. All these resources are tightly coupled 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.

Virtual Storage Platform Architecture

The Virtual Storage Platform offers an entirely new level of scalable enterprise storage, capable of handling the most demanding workloads while maintaining great flexibility. The Virtual Storage Platform offers much higher performance, higher performance scalability, higher reliability and greater flexibility than any storage system on the market today.

The Virtual Storage Platform offers these features:

- The HiStar-E PCI Express Switched Grid acts as the interconnection among front-end directors, back-end directors, data cache adapter boards and virtual storage director boards.
- Data accelerator processors on the front-end directors and back-end directors work with central processor boards called virtual storage directors that manage all I/O by sets of assigned logical devices (LDEVs).
- Dual SAS controllers on back-end director boards contain eight 6Gbps SAS links per board.
- The control memory function resides in global cache and each VSD board contains a local copy with information for its LDEVs. Most control memory accesses are lookups to the local copy.
- Global cache is backed up to solid state drives (SSDs) on the cache boards.
- Each virtual storage director board controls all I/O operations for a discrete group of LDEVs. LDEVs are assigned round-robin across the installed virtual storage director boards as they are created. If necessary, you can manually reassign LDEV ownership to a different virtual storage director.
- Each virtual storage director board executes the code for initiator mode (hosts), external mode (virtualization), back-end director mode, or the copy products send and receive modes. Code execution is done on a per-job basis.
- A Virtual Storage Platform can be scaled from a single-chassis system to a dual-chassis system. Each chassis has a control rack and a logic box.
- Up to 1,280 3.5-inch large form factor (LFF) drives or 2,048 2.5-inch small form factor (SFF) drives can be installed in a dual-chassis system. If you install both LFF and SFF disk containers and drives in a storage system, the limits change based on the configuration you choose.

The Virtual Storage Platform is built as a single-chassis or dual-chassis storage system. Each chassis has one control rack and up to two disk expansion racks. The control rack has the logic box that holds all of the control boards for a chassis and one or two disk containers. The disk expansion racks can hold three disk containers each. Disk containers come in two types: small form factor (up to 128 2.5-inch drives) and large form factor (up to 80 3.5-inch drives). When using two chassis as a single integrated storage system, the two units are cross connected at the grid switch level. The storage system behaves as a single unit, not as a pair of units operating as a cluster.

Figure 2 shows the two types of racks available for a Virtual Storage Platform.

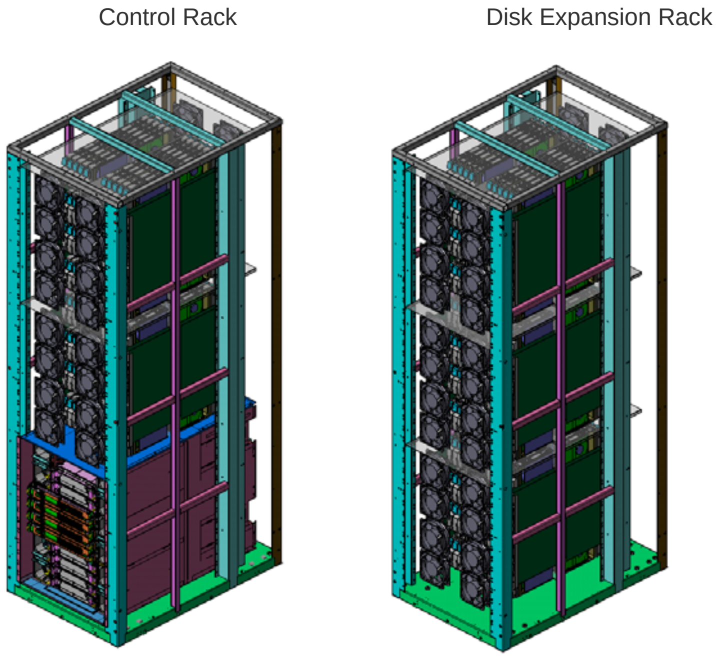

Figure 2

Figure 3 shows the logic boards in a fully populated single-chassis Virtual Storage Platform.

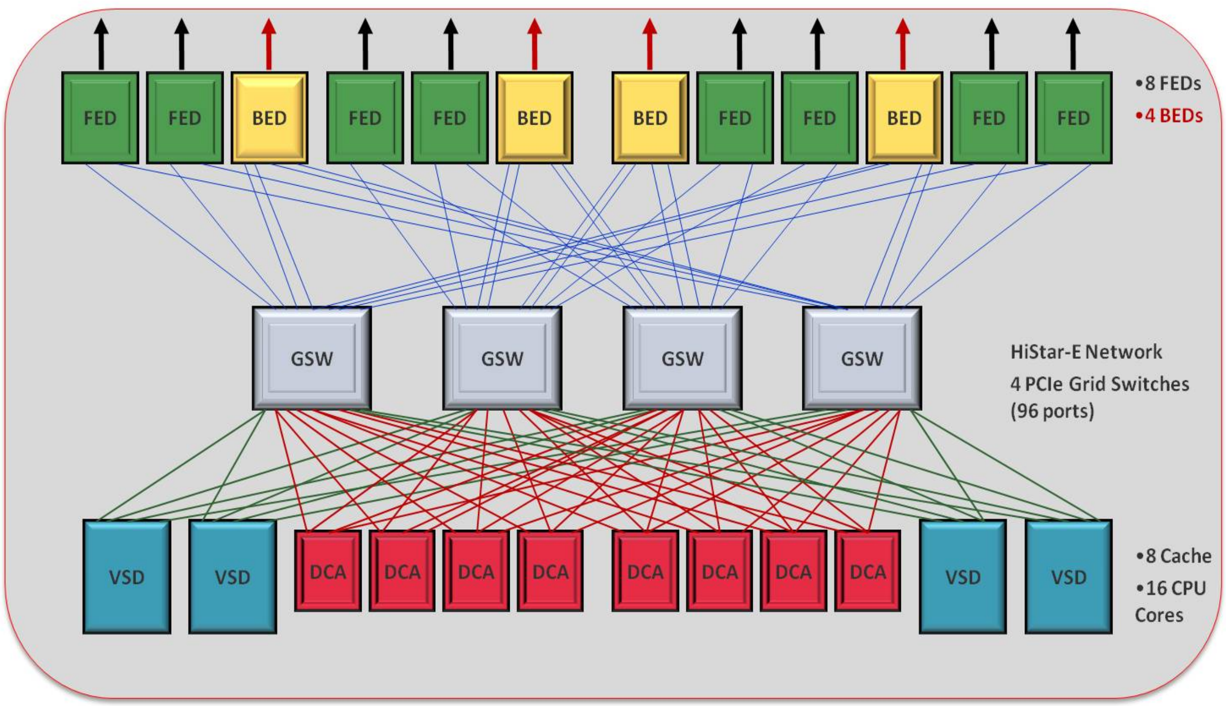

Figure 3

Figure 3 shows the following chassis components (note that a feature is a pair of boards on two separate power domains):

- **GSW** — Grid Switch PCI Express Switch. One or two features (two or four boards) per control unit with 24 2GB/sec HiStar-E ports each can be installed
- **DCA** — Data cache adapter cache memory. One to four features (two, four, six or eight boards) per control unit with up to 32GB of RAM each can be installed
- **VSD** — Virtual storage director processor module. One or two features (two or four boards) per control unit can be installed.
- **FED** — Front-end director host port module. One to four features (two, four, six or eight boards) per control unit of four or eight 8Gbps Fibre Channel ports can be installed.
- **BED** — Back-end director disk controller module. One or two features (two or four boards) per control unit with eight 6Gbps SAS links per board can be installed. If the back-end director options are not installed (available for the single-chassis configuration only), two additional front-end director options can be used in those chassis slots.

3D Scaling Architecture

The Hitachi Virtual Storage Platform allows for optimal infrastructure growth in all dimensions by scaling up, scaling out and scaling deep.

Scale Up

Scale up to meet increasing demands by dynamically adding processors, connectivity and capacity in a single unit, providing the highest performance for both open and mainframe environments.

In the basic single chassis configuration, the number of logic boards, disk containers, and drives is highly scalable. You can start with the minimum set of logic boards,10, and one disk container, then add more boards (up to a total of 28 boards in a single chassis) and disk containers (up to a total of eight disk containers in a single chassis). Disk container types may be intermixed within a chassis.

Scale Out

Scale out to meet multiple demands by dynamically combining multiple units into a single logical system with shared resources, support increased demand in virtualized server environments, and ensure safe multitenancy — that is, the ability to run multiple servers simultaneously without the risk of corruption or modification of data from one server to another — and quality of service through partitioning of cache and ports.

You can double the scalability of the Virtual Storage Platform with a dual-chassis system with up to six racks. The logic box in each chassis is the same, using the same types and numbers of logic boards. Any front-end port can access any back-end RAID group; no division within the storage system exists between the chassis.

A dual-chassis Virtual Storage Platform can manage up to 247PB of total storage capacity.

Table 1 lists the capacity differences between a single-chassis and a dual-chassis Virtual Storage Platform storage system.

Table 1. Virtual Storage Platform Chassis Capacity Comparison

Maximum Capacity	***Single Chassis***	***Dual Chassis***
Data cache	256GB	512GB
Raw cache bandwidth	64GB/s	128GB/s
Solid state drives	128	256
2.5" SFF drives	1,024	2,048
3.5" LFF drives	640	1,280
Logical volumes (LDEVs)	65,280	130,560

Scale Deep

Scale deep to extend storage value by dynamically virtualizing new, existing external storage systems, extend Hitachi Virtual Storage Platform advanced functions to multivendor storage and offload less demanding data to external tiers to optimize the availability of your tier one resources.

The Virtual Storage Platform provides the virtualization mechanisms that allow other storage systems to be attached to some of its front-end director Fibre Channel ports and accessed and managed via hosts that are attached to the host ports on the Virtual Storage Platform. As far as any host is concerned, all virtualized logical units passed through the Virtual storage Platform to the hosts appear to be internal logical units from the Virtual Storage Platform. The front-end ports on the Virtual Storage Platform that attach to the external storage system's front-end ports are operated in external or SCSI initiator mode (attached to servers), rather than the usual SCSI target mode (attached to hosts).

For more information about the Hitachi Virtual Storage Platform, see the Hitachi Data Systems web site.

Storage Area Network Switches

Two Brocade 5340 switches were used in the environment. These switches provide the 8Gb/sec connectivity between the physical servers and the Hitachi Virtual Storage Platform with between 48 to 80 ports total per switch. They are geared towards highly available environments that required a core switch for midsize to large size deployment for a SAN infrastructure. For more information, see the Brocade and Hitachi Data Systems Products & Solutions web site.

Hitachi Dynamic Provisioning Software

On the Virtual Storage Platform, Hitachi Dynamic Provisioning software provides wide striping and thin provisioning functionalities. In the most basic sense, Hitachi Dynamic Provisioning software is similar to the use of a host-based logical volume manager (LVM), but with several additional features available within the Hitachi Virtual Storage Platform and without the need to install software on the host or incur host processing overhead. Hitachi Dynamic Provisioning software provides for one or more pools of wide striping across many RAID groups within a Virtual Storage Platform. One or more Dynamic Provisioning virtual volumes (DP-VOLs) of a user-specified logical size of up to 60TB (with no initial physical space allocated) are created against each pool.

Primarily, you deploy Hitachi Dynamic Provisioning software to avoid the routine issue of hot spots that occur on logical device volumes (LDEVs) from individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. By using many RAID groups as members of a striped Dynamic Provisioning pool underneath the virtual or logical volumes seen by the hosts, a host workload is distributed across many RAID groups, which provides a smoothing effect that dramatically reduces hot spots.

Hitachi Dynamic Provisioning software also carries the side benefit of thin provisioning, where physical space is only assigned from the pool to the DP-VOL as needed using 42MB pages, up to the logical size specified for each DP-VOL. A pool can also be dynamically expanded by adding more capacity or reduced by withdrawing pool capacity. Either operation is performed without disruption or requiring downtime. Upon expansion, a pool can be rebalanced so that the data and workload are wide striped evenly across the current and newly added RAID groups that make up the pool.

Hitachi Dynamic Provisioning software's thin provisioning and wide striping functionalities provide virtual storage capacity to eliminate application service interruptions, reduce costs and simplify administration, as follows:

- Optimizes or "right-sizes" storage performance and capacity based on business or application requirements.
- Supports deferring storage capacity upgrades to align with actual business usage.
- Simplifies the storage administration process.
- Provides performance improvements through automatic optimized wide striping of data across all available disks in a storage pool.
- Eliminates hot spots across the different RAID groups by smoothing the combined workload.
- Significantly improves capacity utilization.

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

Microsoft SharePoint Server 2010

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 such as 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 Server 2010 deployments can range from a standalone environment with a single server for a small amount of users (less than 100) 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. This is made easier through the use of Hyper-V virtual machine templates.

Each tier can be expanded by adding servers that provide a given SharePoint role. 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 into the farm. Servers may also share farm roles which in turn enable a better utilization of the farm servers' resources. Such is the case with using Web Servers to host both the web application as well as central administration and indexing roles.

For more information about Microsoft SharePoint Server 2010, see the Microsoft SharePoint Server web site.

Microsoft SQL Server 2008 R2

Microsoft SQL Server 2008 R2 facilitates the management of any data, any place and any time. Together with the Hitachi Virtual Storage Platform, SQL Server 2008 provides a scalable, high-performance database engine for Microsoft SharePoint Server 2010 and other applications that require the highest levels of availability and security, while reducing the total cost of ownership through enhanced enterprise-class manageability for database deployments

For more information about the features of SQL Server 2008 R2, see the What's New page of SQL Server 2008 R2 Books Online or the Microsoft SQL Server 2008 R2 Product Details web site.

Solution Design

This section provides detailed information on the large scale SharePoint 2010 environment design used for this reference architecture. Although it includes both application and hardware design information, it focuses primarily on the storage design required to build the basic 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, and availability of 20 individual sites hosted under a single site collection.

The architecture is built with Hyper-V as the hypervisor for all of the SharePoint farm server infrastructure and Hitachi Dynamic Provisioning software 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 SharePoint Server deployments.

Figure 4 shows the SharePoint infrastructure from a server, storage, and network perspective.

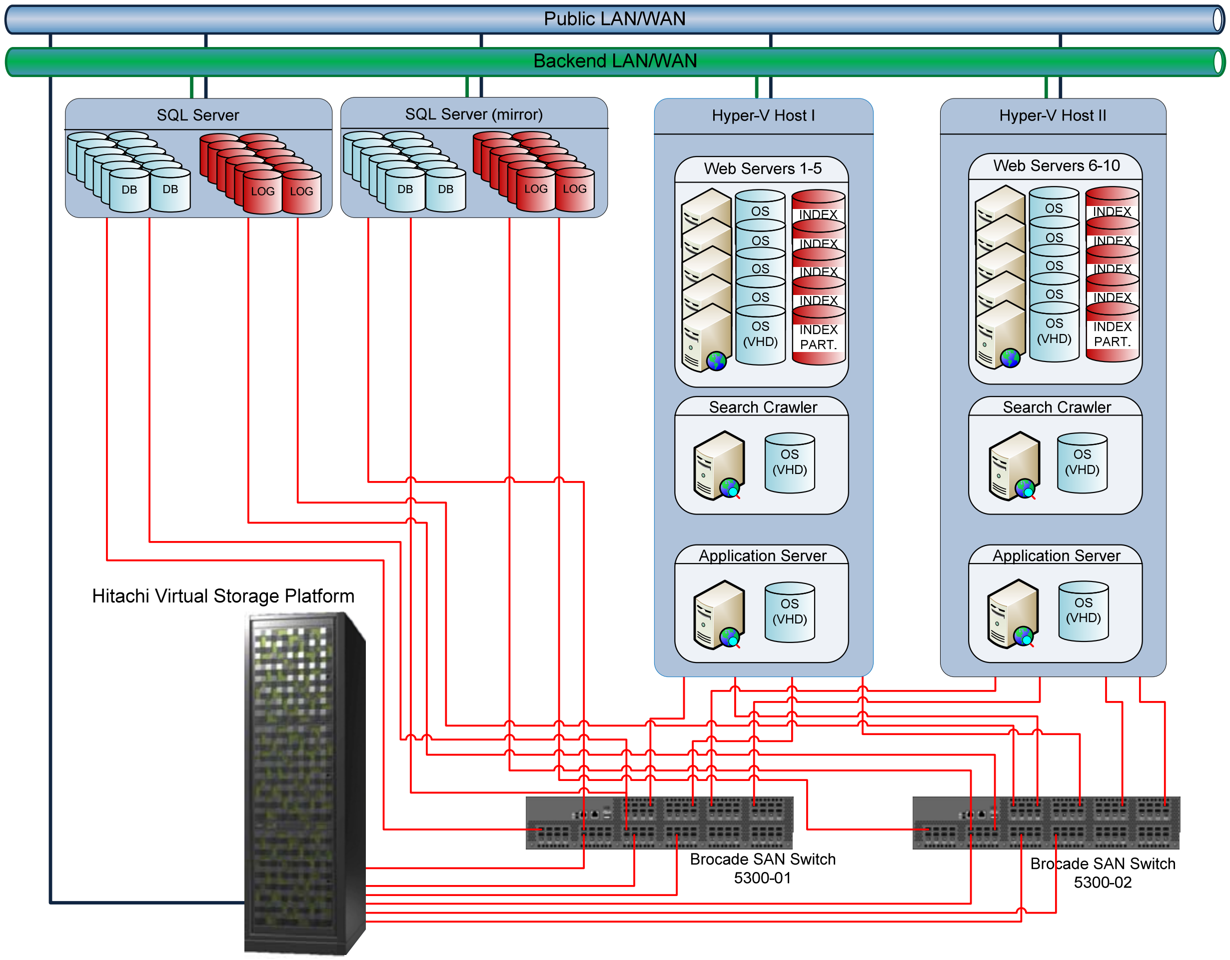

Figure 4

Storage Architecture

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

This solution uses a 6TB reference architecture that stores all SharePoint databases and index volumes with corresponding database files, transaction logs and `tempdb` data. Using Hitachi Dynamic Provisioning software simplified deployment because it enabled the quick creation of the environment volumes after the Dynamic Provisioning pools were created.

Architecture Capacity Planning

The capacity and performance requirements for a SharePoint environment are dictated by many variables such as the number of documents and their size, the type of content being saved, and number of users that will be using the environment. The capacity required by your architecture is primarily dictated by the amount of content that will be saved on the SharePoint environment and the web site structure of your environment. For optimum performance, Hitachi and Microsoft recommend using RAID-1+0 for all of the volumes in a SharePoint environment. However, with Hitachi Dynamic Provisioning software, you can also consider using RAID-5 for the SharePoint environment if appropriate testing is done to ensure that it meets your performance requirements. Keep the following key best practices in mind:

- Limit each content database to 200GB, unless the data remains static for archiving purposes, in which case Microsoft supports 1TB content databases.
- If you plan to use SharePoint Server 2010 granular backup tools, limit the size of each site collection to 100GB.
- The search administration database is 10GB in size
- The minimum size of the crawl database is 4.6 percent of the sum of all content database sizes.
- The minimum size of the search property database is 1.5 percent of the sum of all content database sizes.
- The minimum size of the configuration database is 2GB. Note that 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 all content databases multiplied by four; this accounts for index repartitioning.
- In most cases, the minimum size of transaction logs for each 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 databases that exceed 100GB, provide a fixed number in megabytes for the autogrowth value instead of a percentage. This reduces how often these file sizes must be increased and therefore minimizes the possible performance overhead on both the server and the storage system caused by the increase in the file sizes. Set the MAXSIZE setting on each database file to a value that matches the capacity of the volume being used. You can add volumes and files to the database later if needed.

You might also need to consider additional storage requirements depending on the features you plan to install in your SharePoint environment. This reference architecture guide 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 SharePoint deployment, see the Microsoft TechNet article "Storage and SQL Server capacity planning and configuration."

Architecture IOPS Planning

When determining your IOPS requirements, keep the following guidelines from Microsoft in mind:

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

In addition, keep the following key best practices in mind:

- Create dedicated volumes for each of the databases used in SharePoint in order to allow the queue depth to be applied. The queue depth is normally applied on a LUN basis on the host side.
- 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. In addition, turn on the MAXSIZE setting for each of the database files 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 all files 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 SQL Server Developer Center article "Optimizing tempdb Performance."

No precise method exists to establish content database IOPS requirements, and like any other SQL Server environment, best practice is to first deploy a test environment to establish whether your content databases can support your IOPS requirements.

Environment Storage Architecture

The entire tested architecture used 300GB 10K RPM SAS drives configured as RAID-1+0 (2D+2D) groups assigned to two dedicated Dynamic Provisioning pools. One pool was used to host all of the database and `tempdb` volumes required for the environment as well as the index volumes assigned to each individual Web server. The second pool was used to host all of the transaction log volumes required for the environment.

This reference architecture includes a pool shared within the infrastructure to host the boot OS for the virtual machines. Note that although the reference architecture used database mirroring at the SQL Server level for high availability, those volumes are not documented in the following tables. However, they used the same number of pools and RAID group configurations as the production environment. Table 2 lists the size of disks and RAID group configuration for each pool.

Table 2. Dynamic Provisioning Pools

Pool Name (ID)	*RAID Group Configuration*	*Drive Type*	*Number of RAID Groups*	*Pool Capacity (TB)*
HDP-MSFT-SP-DB00 (30)	RAID-1+0 (2D+2D)	300GB 10K RPM SAS	11	5.70
HDP-MSFT-SP-LOG00 (31)	RAID-1+0 (2D+2D)	300GB 10K RPM SAS	2	1.00
HDP-MSFT-OS00 (00)	RAID-5 (3D+1P)	2TB 7.5K RPM SATA	1	5.37

Although 11 RAID groups were used for the HDP-MSFT-SP-DB00 pool, fewer RAID groups can be used initially due to the thin provisioning capability of Hitachi Dynamic Provisioning software. However, using the amount needed for the deployment with no oversubscription can provide better performance due to the wide striping capability of Hitachi Dynamic Provisioning software.

Table 3 lists the volumes provisioned for the configuration database out of each Dynamic Provisioning pool.

Table 3. Configuration Database Volumes

Pool Name (ID)	*LDEV*	*Size (GB)*	*Purpose*	*Storage Port*
HDP-MSFT-SP-DB00 (30)	30:00	2	Database	1D/2D
HDP-MSFT-SP-LOG00 (31)	31:00	1	Log	3D/4D

Table 4 lists the volumes provisioned for the central administration database out of each Dynamic Provisioning pool.

Table 4. Central Administration Database Volumes

Pool Name (ID)	*LDEV*	*Size (GB)*	*Purpose*	*Storage Port*
HDP-MSFT-SP-DB00 (30)	30:01	1	Database	3D/4D
HDP-MSFT-SP-LOG00 (31)	31:01	1	Log	1D/2D

Table 5 lists the volumes provisioned for the content databases out of each Dynamic Provisioning pool.

Table 5. Content Databases Volumes

Pool Name (ID)	*LDEVs*	*Size (GB)*	*Purpose*	*Storage Port*
HDP-MSFT-SP-DB00 (30)	30:02, 30:04, 30:06, 30:08, 30:0A, 30:0C, 30:0E, 30:10, 30:12, 30:14	200	Databases 00, 02, 04, 06, 08, 10, 12, 14, 16, 18	1D/2D
HDP-MSFT-SP-DB00 (30)	30:03, 30:05, 30:07, 30:09, 30:0B, 30:0D, 30:0F, 30:11, 30:13, 30:15	200	Databases 01, 03, 05, 07, 09, 11, 13, 15, 17, 19	3D/4D
HDP-MSFT-SP-LOG00 (31)	31:02, 31:04, 31:06, 31:08, 31:0A, 31:0C, 31:0E, 31:10, 31:12, 31:14	40	Logs 00, 02, 04, 06, 08, 10, 12, 14, 16, 18	3D/4D
HDP-MSFT-SP-LOG00 (31)	31:03, 31:05, 31:07, 31:09, 31:0B, 31:0D, 31:0F, 31:11, 31:13, 31:15	40	Logs 01, 03, 05, 07, 09, 11, 13, 15, 17, 19	1D/2D

Table 6 lists the volumes provisioned for the search administration database out of each Dynamic Provisioning pool.

Table 6. Search Administration Database Volumes

Pool Name (ID)	LDEV	Size (GB)	Purpose	Storage Port
HDP-MSFT-SP-DB00 (30)	30:16	10	Database	3D/4D
HDP-MSFT-SP-LOG00 (31)	31:16	2	Log	1D/2D

Table 7 lists the volumes provisioned for the crawl database out of each Dynamic Provisioning pool.

Table 7. Crawl Database Volumes

Pool Name (ID)	LDEV	Size (GB)	Purpose	Storage Port
HDP-MSFT-SP-DB00 (30)	30:17	185	Database	1D/2D
HDP-MSFT-SP-LOG00 (31)	31:17	36	Log	3D/4D

Table 8 lists the volumes provisioned for the search property database out of each Dynamic Provisioning pool.

Table 8. Search Property Database Volumes

Pool Name (ID)	LDEV	Size (GB)	Purpose	Storage Port
HDP-MSFT-SP-DB00 (30)	30:18	60	Database	3D/4D
HDP-MSFT-SP-LOG00 (31)	31:18	15	Log	1D/2D

Table 9 lists the volumes provisioned for the `tempdb` file out of each Dynamic Provisioning pool.

Table 9. `tempdb` Volumes

Pool Name (ID)	LDEV	Size (GB)	Purpose	Storage Port
HDP-MSFT-SP-DB00 (30)	30:30	320	Database	3D/4D
HDP-MSFT-SP-DB00 (30)	30:31	320	Database and log	1D/2D

Table 10 lists the volumes provisioned for the index volumes out of each Dynamic Provisioning pool.

Table 10. Index Volumes

Pool Name (ID)	LDEV	Size (GB)	Purpose	Storage Port
HDP-MSFT-SP-DB00 (30)	30:35	60	Web server 1	3D/4D
HDP-MSFT-SP-DB00 (30)	30:36	60	Web server 2	1D/2D
HDP-MSFT-SP-DB00 (30)	30:37	60	Web server 3	3D/4D
HDP-MSFT-SP-DB00 (30)	30:38	60	Web server 4	1D/2D
HDP-MSFT-SP-DB00 (30)	30:39	60	Web server 5	3D/4D
HDP-MSFT-SP-DB00 (30)	30:3A	60	Web server 6	5D/6D
HDP-MSFT-SP-DB00 (30)	30:3B	60	Web server 7	7D/8D
HDP-MSFT-SP-DB00 (30)	30:3C	60	Web server 8	5D/6D
HDP-MSFT-SP-DB00 (30)	30:3D	60	Web server 9	7D/8D
HDP-MSFT-SP-DB00 (30)	30:3E	60	Web server 10	5D/6D

Table 11 lists the volumes provisioned for the boot OS volumes out of each Hitachi Dynamic Provisioning Pool.

Table 11. Boot OS Volumes

Pool Name (ID)	LDEV	Size (GB)	Purpose	Storage Port
HDP-MSFT-OS00 (00)	03:00	600	Hyper-V virtual machine VHD repository 1	7B/8B
HDP-MSFT-OS00 (00)	03:01	600	Hyper-V virtual machine VHD repository 2	7B/8B

Although the SQL Server databases are the most important aspect of the overall SharePoint environment volume requirements, they are only one aspect of the overall environment. SharePoint content index utilized by the search functionality is maintained as flat files on NTFS volumes. The content index volumes in this reference architecture utilized the same SAN environment and Dynamic Provisioning pool used for the environment database volumes.

Architecture Scaling

You can easily scale the storage in your environment by provisioning additional content databases volumes (database and log) out of existing Dynamic Provisioning pools. However, be careful when oversubscribing the pool capacity with the thin provisioning capability of Dynamic Provisioning pools. You must monitor the capacity utilization of each pool and add pool volumes when necessary.

In this reference architecture, each additional RAID-1+0 (2D+2D) group provides 530GB of total capacity to the pool. This means that you can add up to two 200GB content database volumes to the pool dedicated to database volumes (HDP-MSFT-SP-DB00). Note that with the addition of content databases to your environment, other database sizes must also be recalculated to continue to adhere to Microsoft's capacity best practices. Due to its wide striping capability, Hitachi Dynamic Provisioning software can reduce hot spots that can degrade performance. This means that although you only use one RAID group for two content databases, you can spread your I/O workload across all of the spindles that are part of the Dynamic Provisioning pool. Use site collection quota management to monitor and prevent excess storage requirements and plan your growth appropriately.

For more information about calculating your capacity and IOPS requirements, see the "Architecture Capacity Planning" section of this paper or the Microsoft TechNet article "Storage and SQL Server capacity planning and configuration."

Host Considerations

To ensure optimal performance from your storage system, it is important to understand disk alignment, NTFS allocation unit size, path redundancy, HBA placement and queue depth during the planning process.

Disk Alignment

You must set both the disk alignment offset and the NTFS allocation unit size when LUNs are partitioned and formatted at the operating system level before you create Microsoft SQL Server database files. Windows Server 2008 eliminated the need to align the partition due to its use of 1024K as the default partition offset. It is important to remember that this is a destructive process that must be performed before database files are created to prevent data loss.

NTFS Allocation Unit Size

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

Host Bus Adapter Drivers

When choosing HBA drivers, ensure that you are using the current recommended drivers for the 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 the Interoperability Information section 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 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.

For more information, see the Microsoft TechNet article "Monitoring and maintaining SharePoint Server 2010."

SharePoint Architecture

This solution uses a three-tier architecture that provides the basic functionalities for a large collaboration environment by using ten web servers with the network load balance feature enabled and search crawl servers that provide the crawling functionality to the farm. The web servers in the environment also served as the search index servers by having a distributed index architecture in accordance with Microsoft's recommendations. SharePoint farms allow scaling by adding servers to the environment as the number of users or the amount of content increases. Additions such as servers within each tier can be made at any time simply by connecting the new hardware or deploying new virtual machines and running the SharePoint configuration wizard. Note that although a second SQL server is included in this reference architecture and used as the server hosting the mirrored databases for the environment, the volumes used by this server are not documented in this guide. For more information about setup and requirements for SQL Server database mirroring, see the Database Mirroring Setup Overview TechNet page.

Figure 5 shows the applications that each virtual machine hosts.

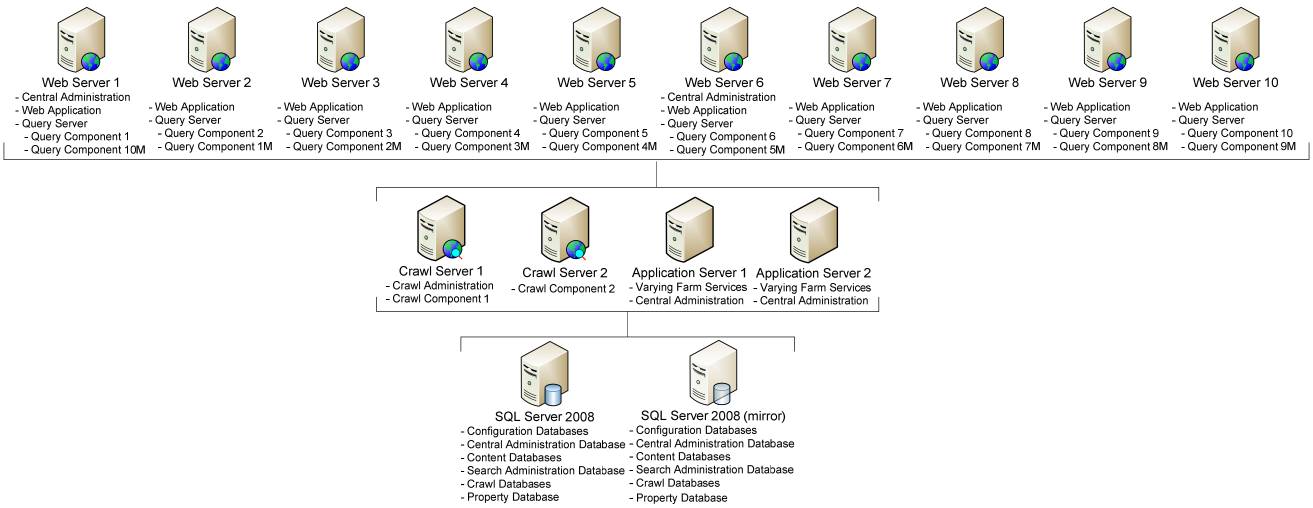

Figure 5

The number of web servers that your environment requires depends primarily on the amount of strain 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 are the CPU, memory and network utilization rates in comparison to their limits. For example, if you find that 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, consider adding web servers.

SQL Server Architecture

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 each file type's workload, 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.

Both Microsoft and Hitachi Data Systems recommend using SQL Server 2008 R2 Enterprise edition for your SharePoint Server 2010 environment to take advantage of the additional performance, availability, security and management capabilities that it has in comparison to the Standard edition.

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, 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. You can 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. Both Microsoft and Hitachi recommend provisioning separate dedicated volumes for database and transaction log files. If you design your storage properly, the Hitachi Virtual Storage Platform storage system's optimized caching operations and the use of Hitachi Dynamic Provisioning software ensure that logs can be written without delay.

One of the features built into the Virtual Storage Platform is the ability to optimize physical I/O based on recognition of I/O patterns. The Virtual Storage Platform storage system's 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 Virtual Storage Platform 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, Hitachi's testing shows that it can reside on the same RAID group 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. You must restart SQL Server when you add a new file.

Database Performance

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.

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. For example, in this solution, scaling the search database as content growth increases was addressed by adding additional volumes to the dynamic provisioning pool when such became necessary either from a capacity or performance perspective, and the creation of 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 physical server being used for the 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. The physical servers utilized in this solution were used as the Microsoft Hyper-V hosts for the virtual machines in the SharePoint environment. Two additional physical servers were utilized for the deployment of SQL server due to the performance needs of the environment. The physical servers were chosen due to their hypervisor support and hardware capabilities to provide processor, memory, and overall internal I/O bandwidth required to support a large enterprise level SharePoint farm. Two physical servers were utilized in this solution in order to provide the resources needed for the SharePoint farm infrastructure. Whether you implement your SharePoint environment on physical or virtual servers, you must keep the physical server architecture in mind.

The reference architecture for this paper was deployed on a virtual environment using Microsoft Windows Hyper-V as the hypervisor. Table 12 lists the physical servers used in this reference architecture.

Table 12. Physical Server Configuration

Server	*Description*	*Firmware Level*	*Quantity*	*Role*	*Operating System*
Hyper-V host servers	4 x Quad-Core AMD Opteron processor 1.9GHz, 128GB RAM. Equipped with 2 x Emulex LPe11002 4GB HBAs.	BIOS firmware 4.1.1	2	Windows Server Hyper-V host	Windows Server 2008 R2 Data Center Edition
SQL servers	4 x Quad-Core AMD Opteron processor 1.9GHz, 64GB RAM. Equipped with 2 x Emulex LPe11002 4GB HBAs.	BIOS firmware 4.1.1	2	SQL Server 2008 R2	Windows Server 2008 R2 Enterprise Edition

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

Table 13. Virtual Server Configuration

Name	Role	Server	Quantity	Drives	Details
RAID700-MSFT-SP-WS	Web servers	Hyper-V host 1	5	127GB VHD – operating system 50GB VHD – SharePoint Server 2010 files	Windows 2008 R2 Enterprise Edition, 2 VCPUs, 8GB RAM
RAID700-MSFT-SP-AS	Crawl server	Hyper-V host 1	1	127GB VHD – operating system 50GB VHD – SharePoint Server 2010 files	Windows 2008 R2 Enterprise Edition, 2 VCPUs, 12GB RAM
RAID700-MSFT-SP-APP1	Application server	Hyper-V host 1	1	127GB VHD – operating system	Windows Server 2008 R2 Enterprise Edition, 2 VCPUs, 8GB RAM.
RAID700-MSFT-SP-WS	Web servers	Hyper-V host 2	5	127GB VHD – operating system 50GB VHD – SharePoint Server 2010 files	Windows 2008 R2 Enterprise Edition, 2 VCPUs, 8GB RAM
RAID700-MSFT-SP-AS	Crawl server	Hyper-V host 2	1	127GB VHD – operating system 50GB VHD – SharePoint Server 2010 files	Windows 2008 R2 Enterprise Edition, 2 VCPUs, 12GB RAM
RAID700-MSFT-SP-APP2	Application server	Hyper-V host 2	1	127GB VHD – operating system	Windows Server 2008 R2 Enterprise Edition, 2 VCPUs, 8 GB RAM.

The virtual servers used in the SharePoint infrastructure were distributed between two physical hosts for both high availability and resource utilization. Ten web servers were created and distributed evenly between the two physical servers. Each web server also contained one index partition and one mirror index partition from another web server for high availability and indexing performance. Each physical server also hosted one virtual machine with a search crawl component for the environment. A total of two SQL Server virtual machines were used for the environment in a mirrored configuration for high availability purposes.

Hyper-V virtual machines can access LUNs in two ways: using pass-through disks or virtual hard disks (VHDs) created on an LUN. Pass-through disks enable the LUNs to be available to a virtual machine as if they were directly connected to it. For the virtual machines operating systems, the LUNs are presented to the physical host and used as containers to house the virtual machine OS VHD file or additional data files. This reference architecture uses pass-through disks for all of the volumes, including the database, `tempdb` and log files for each of the databases, assigned to the virtual machines.

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 SharePoint environment were mapped to multiple ports on the Hitachi Virtual Storage Platform. The LDEVs were assigned to each port in a way that ensures that the host had two paths to the storage system for both high availability as well as load balancing. Additionally, the database and transaction volumes (LDEVs) for each database were mapped to alternate ports. This ensures balancing so that the physical database has its own primary path prioritization. For the volumes provisioned for the boot OS environment, two ports that were shared amongst other hosts were used.

The environment used two Brocade 5340 switches to provide scalability and high availability to the environment. One port of each of the physical host's HBA was connected to a switch to provide high availability and the four ports from the Hitachi Virtual Storage Platform were also divided between the two switches with high availability in mind.

Table 14 lists the zoning details for the SAN.

Table 14. SAN Switch Architecture

Server	HBA Ports	Switch Zone	Storage Port	Switch
Hyper-V host 1	HBA1-1	HYPERV1_HBA1_1_VSP-101_1D_7B	1D, 7B	5300-01
Hyper-V host 1	HBA1-2	HYPERV1_HBA1_2_VSP-101_2D	2D	5300-02
Hyper-V host 1	HBA2-1	HYPERV1_HBA2_1_VSP-101_3D_8B	3D, 8B	5300-01
Hyper-V host 1	HBA2-2	HYPERV1_HBA2_2_VSP-101_4D	4D	5300-02
Hyper-V host 2	HBA1-1	HYPERV2_HBA1_1_VSP-101_5D_7B	5D, 7B	5300-01
Hyper-V host 2	HBA1-2	HYPERV2_HBA1_2_VSP-101_6D	6D	5300-02
Hyper-V host 2	HBA2-1	HYPERV2_HBA2_1_VSP-101_7D_8B	7D, 8B	5300-01
Hyper-V host 2	HBA2-2	HYPERV2_HBA2_2_VSP-101_8D	8D	5300-02

Server	***HBA Ports***	***Switch Zone***	***Storage Port***	***Switch***
SQL Server (production)	HBA1-1	SQLSERVER_P_HBA1_1_VSP-101_1D_7B	1D, 7B	5300-01
SQL Server (production)	HBA1-2	SQLSERVER_P_HBA1_2_VSP-101_2D	2D	5300-02
SQL Server (production)	HBA2-1	SQLSERVER_P_HBA2_1_VSP-101_3D_8B	3D, 8B	5300-01
SQL Server (production)	HBA2-2	SQLSERVER_P_HBA2_2_VSP-101_4D	4D	5300-02
SQL Server (mirror)	HBA1-1	SQLSERVER_M_HBA1_1_VSP-101_1D_7B	1D, 7B	5300-01
SQL Server (mirror)	HBA1-2	SQLSERVER_M_HBA1_2_VSP-101_2D	2D	5300-02
SQL Server (mirror)	HBA2-1	SQLSERVER_M_HBA2_1_VSP-101_3D_8B	3D, 8B	5300-01
SQL Server (mirror)	HBA2-2	SQLSERVER_M_HBA2_2_VSP-101_4D	4D	5300-02

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

- 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 dual SAN fabrics, multiple HBAs and host-based multipathing software when using 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 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 round robin as the load-balancing setting for either Hitachi Dynamic Link Manager software or Windows Server 2008 R2 native MPIO.

Physical and Virtual Network Architecture

As per Microsoft’s best practice, the Hyper-V hosts in this reference architecture have separate 1Gpbs physical network adapters to be used independently for different types of data traffic. One of the adapters is used primarily for Hyper-V host data traffic for management purposes, and the other two are dedicated for virtual machine traffic. This configuration provides improved performance and adapter redundancy. Each Hyper-V host is configured with two virtual LANs that provide the network connections to the virtual machines.

Each virtual machine was configured with two network adapters. One of the adapters provides a dedicated link between the servers as well as for management purposes; the second provides a dedicated channel between the users and the environment servers (primarily the web servers). The Network Load Balancing feature was installed on all web servers and configured to use the network interface dedicated for user connections to the environment.

Figure 6 provides a detailed view of how the network was configured for the environment.

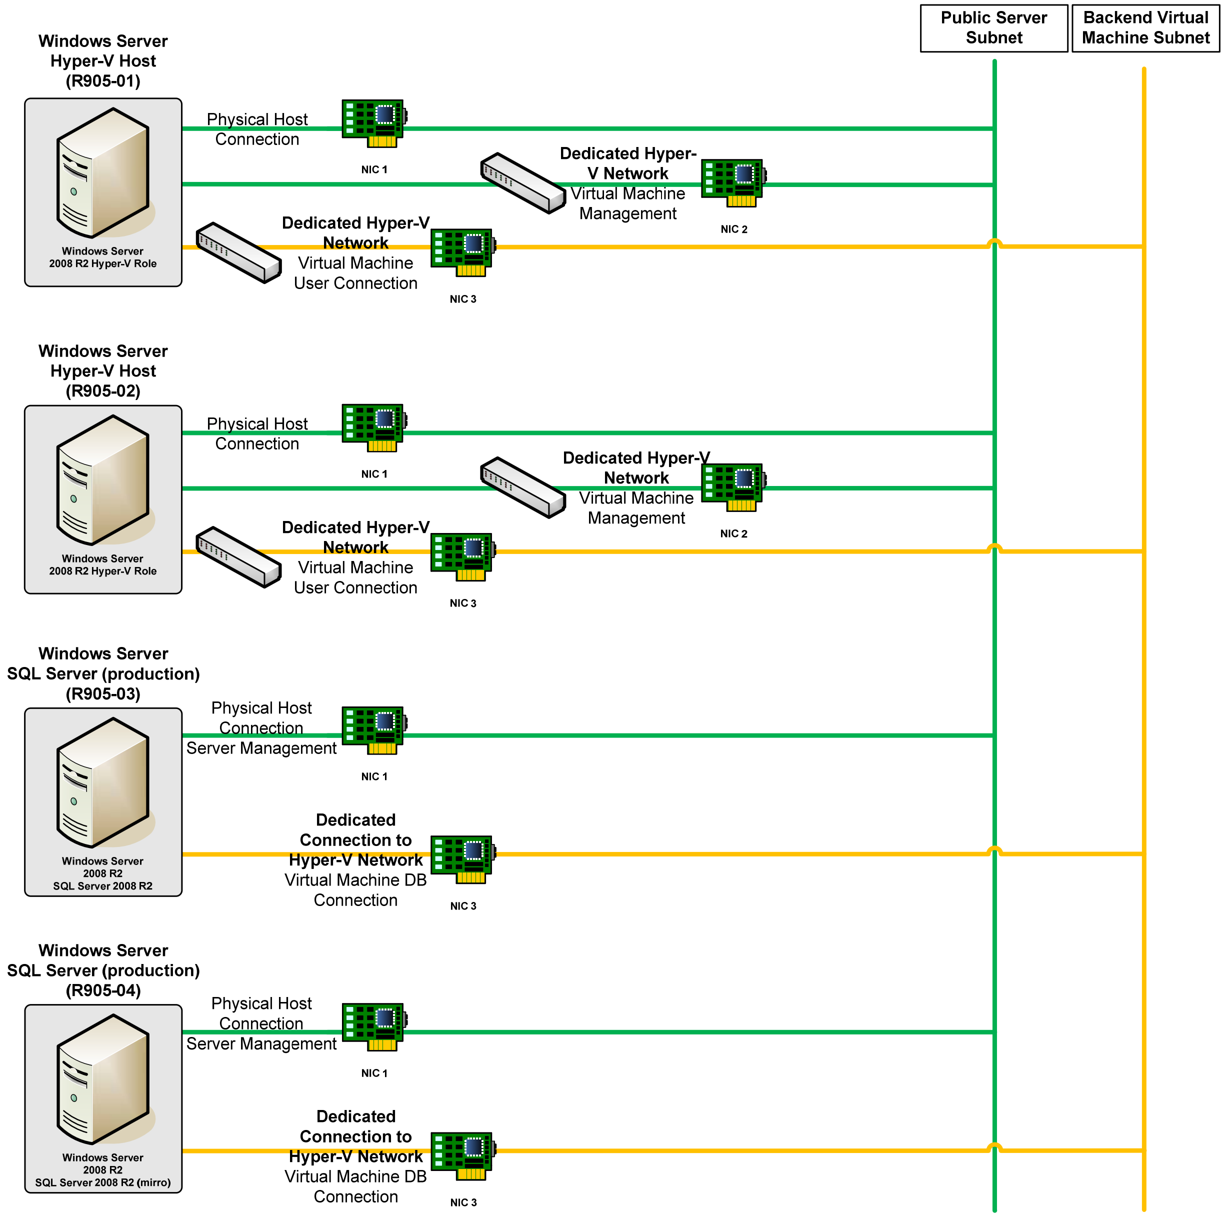

Figure 6

Engineering Validation

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

Test Methodology

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

Table 15. User Load Type Mix

Load Type	***Percentage***
Browsing	80
Upload	10
Search	5
Open	5

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

A total of 200,000 users were created and given contributor permission to the environment's sites. A 1 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 KBs to a couple MBs, however the 100MB default limitation was maintained for upload size. The content databases were first evenly populated with a total of a little more than 1TB of data files of different types. Approximately 50GB worth of data was populated to each site prior to the test execution in order to provide both search and open load types to the environment. For the testing execution, the data files were randomly uploaded to all of the collaboration sites provisioned by different 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/sec. Table 16 lists acceptable page response times that were used for the testing of the environment. All volume response times were within Microsoft's guidelines.

Table 16. Acceptable Response Times

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

Note that 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 different size documents were used (from 10KB to 5MB), 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 greatly 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 they 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 17. Response Times Results

Load Type	Response Time (seconds)
Browsing	<3
Upload	Variable ≤6
Search	<4
Open	Variable ≤3

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

Table 18. Average and Maximum CPU Utilization Rates

Type	***Average CPU Utilization Rate***	***Maximum CPU Utilization Rate***
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 provides high availability and simplified storage administration for enterprise SharePoint deployments. This reference architecture simplifies planning and deploying Microsoft SharePoint Server by leveraging the capabilities of the Hitachi Virtual Storage Platform and Hitachi Dynamic Provisioning software.

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 Corporate Resources web site. Click the **Product Demos** tab for a list of available recorded demonstrations.

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

Corporate Headquarters
750 Central Expressway,
Santa Clara, California 95050-2627 USA
www.hds.com

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