HITACHI
Inspire the Next

Deploying Microsoft SharePoint Server 2010 on the Hitachi Virtual Storage Platform

Implementation Guide

By Eduardo Freitas

January 2011

Microsoft GOLD CERTIFIED Partner

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

Implementation Guide

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 the deployment a 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, Windows, SQL and SharePoint implementation concepts.

Solution Overview

This white paper describes how to deploy a SharePoint Server 2010 environment that supports 200,000 users tested with 20 site collections (using 20 content databases). 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, The test environment included a Windows Domain Controller supplying Active Directory and DNS services. that was already part of the existing lab infrastructure. 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 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.

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

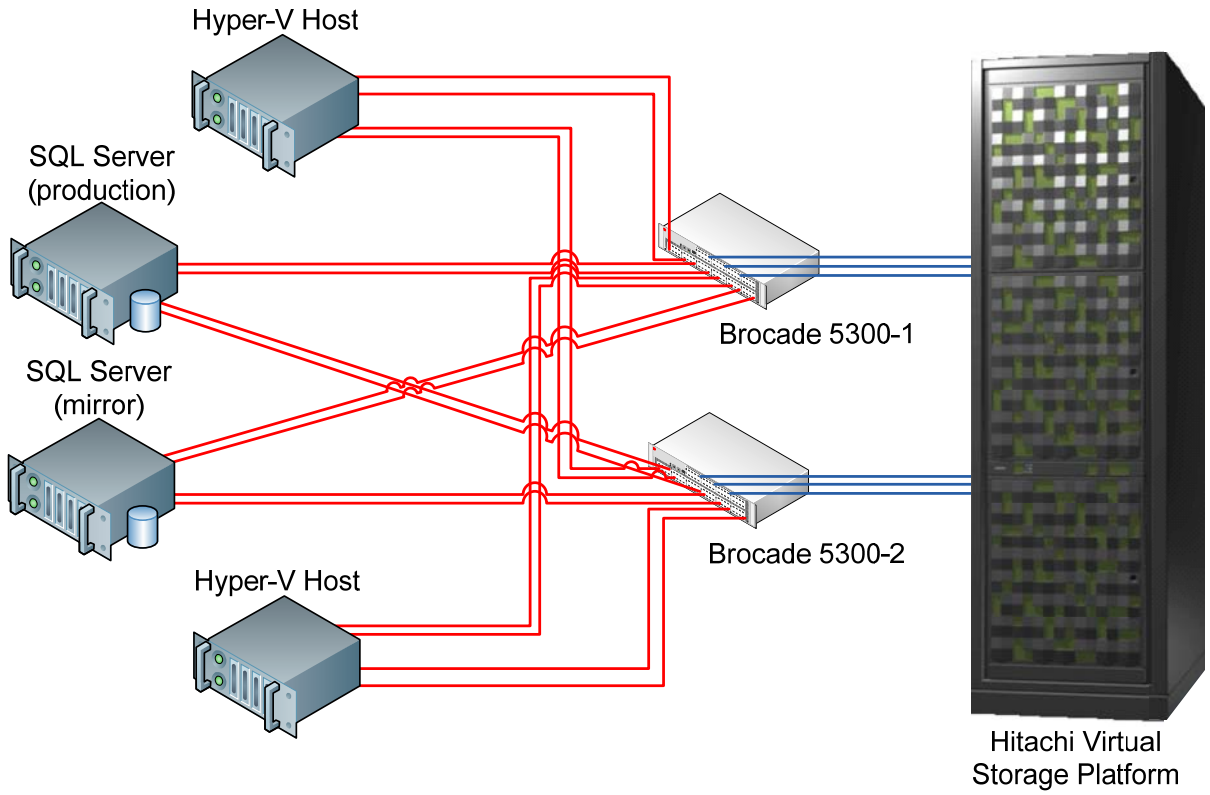

Figure 1

The architecture is build 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 2 shows the SharePoint infrastructure from a server perspective.

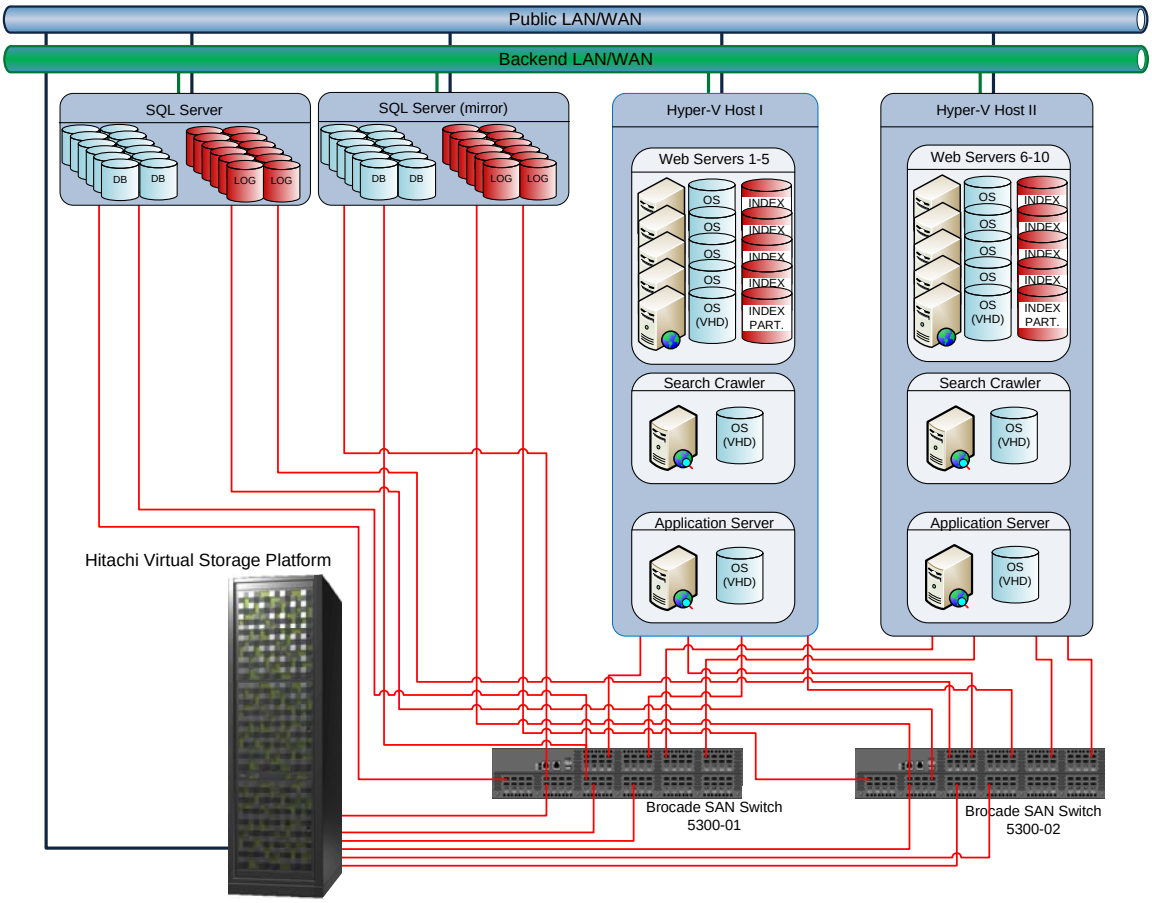

Figure 2

This implementation guide enables storage and database administrators to successfully deploy SharePoint environments by describing the key components necessary for the environment and the tasks required to successfully deploy it.

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 implementation guide 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.

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 Windows Hyper-V

Microsoft Windows Hyper-V is a hypervisor-based virtualization technology that is integrated into Windows Server 2008 x64 and Windows Server 2008 R2 versions of the operating system. It allows for the reduction of hardware footprints and capital expenses through server consolidation. This is accomplished by consolidating multiple physical servers that are hosting SQL Server instances into a single Hyper-V server.

Additional options are available with Hyper-V, such as quick and live migration, which provide high availability for SQL Server virtual machines. One of the requirements for making the SQL Server virtual machines highly available is that they must be hosted in a Hyper-V failover cluster configuration.

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.

Tested Solution Components

The following sections detail the hardware and software components that were deployed in the Hitachi Data Systems lab to support the solution described in this deployment guide.

Hardware Components

Table 1 lists the detailed information about the hardware components used.

Hardware	*Description*	*Firmware Level*	*Quantity*
Hitachi Virtual Storage Platform storage system	Single frame 6 X 8GB Fibre Channel ports used 128GB cache memory 56 X 300GB, 10K RPM SAS disks	70-00-50-00/00	1
Brocade 5300 switch	SAN switch with 8Gb Fibre Channel ports	FOS 6.4.0b	2
Hyper-V Hosts	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
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

Storage Area Network Components

The storage area network (SAN) configuration for this solution uses two Fibre Channel switches for high availability. Redundant paths are configured from the Hyper-V host for the SQL Server 2008 R2 databases, tempdb and transaction logs to the Virtual Storage Platform. For the SQL Server virtual machines the VHDs for the guest OS are configured with two redundant paths to the Virtual Storage Platform. The server has two dual port host bus adapters (HBAs) installed for high availability purposes.

The Microsoft MPIO software is used for multipathing, employing the round-robin multipathing policy. Microsoft MPIO software's round-robin load balancing algorithm automatically selects a path by rotating through all available paths, thus balancing the load across all available paths and optimizing IOPS and response time.

Software Components

Table 2 lists the detailed information about the software components used.

Table 2. Software Components

Software	*Version*
Hitachi Storage Navigator	Dependent on microcode version
Hitachi Dynamic Provisioning	Dependent on microcode version
Windows Server 2008 R2	Datacenter edition for Hyper-V hosts Enterprise edition for SQL Server hosts and VMs
Microsoft SQL Server 2008 R2	Enterprise edition
Microsoft SharePoint 2010	Enterprise edition

Table 3 lists the virtual servers used in this reference architecture.

Table 3. 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.

Name	Role	Server	Quantity	Drives	Details
RAID700-MSFT-SP-WS	Web servers	Hyper-V host 2t	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, 8GB RAM.

Solution Implementation

To deploy this SharePoint Server 2010 solution, follow these high-level steps:

1. Configure fabric switch zones.
2. Configure the Virtual Storage Platform.
3. Configure servers and VMs.
4. Install and configure SQL Server 2008 R2.
5. Install and configure SharePoint Server 2010.
6. Replicate active databases to mirror SQL Server.

These are general tasks that need to be completed for a successful deployment. Your checklist might vary based on your environment. Details about each of these steps are included in the following sections.

For more information about each of these high-level tasks, see the documentation provided by Hitachi Data Systems and Microsoft:

- Hitachi Storage Navigator online help
- Hitachi Virtual Storage Platform Dynamic Provisioning Software user's guide
- Microsoft TechNet articles:
  - "SQL Server and storage (SharePoint Server 2010)"
  - "Deployment for SharePoint 2010"

Configure Fabric Switch Zones

Configure zones on your fabric switches according to the manufacturer's guidelines; in addition, follow these best practices:

- Use World Wide Port Name (WWPN) identification for all zoning configuration.
- Connect a minimum of two HBAs per server for multipath high availability.
- Disable all unused switch ports to increase security.
- Configure ports for point-to-point topology.
- Set ports to a specific speed; do not use the auto negotiate setting.
- Use single initiator zoning.

Table 4 lists the zoning details for the SAN.

Table 4. SAN Switch Architecture

Server	*HBA Ports*	*Switch Zone*	*Storage Port*	*Switch*
Hyper-V Host 1	HBA1-1	HYPERV1_HBA1_1_VSP-101_1D_8B	1D, 8B	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_8B	5D, 8B	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
SQL Server (production)	HBA1-1	SQLSERVER_P_HBA1_1_VSP-101_1D_8B	1D, 8B	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

Server	HBA Ports	Switch Zone	Storage Port	Switch
SQL Server (mirror)	HBA1-1	SQLSERVER_M_HBA1_1_VSP-101_1D_8B	1D, 8B	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

Configure the Virtual Storage Platform

This guide uses two Dynamic Provisioning pools, each created from multiple LDEVs from RAID-1+0 (2D+2D) groups on the Virtual Storage Platform. These procedures assume that RAID groups and LDEVs are already created, which in the RAID groups case is done when Hitachi installs your Virtual Storage Platform. Two pools are used for the environment, one for the databases (including tempdb) and crawl volumes, and the other for the log volumes.

Hitachi Dynamic Provisioning software is used to avoid the routine issue of hot spots that occur on logical devices (LDEVs) from individual RAID groups when the host workload exceeds the IOPS or throughput capacity of that RAID group. By using many LDEVs from different 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. This provides a smoothing effect that dramatically reduces hot spots such as the ones that may occur during full crawling of the environment.

For this solution, Hitachi Data Systems used Hitachi Dynamic Provisioning software to provision the storage for the environment. The environment includes a pool shared within the infrastructure to host the boot OS for the virtual machines. Table 2 lists the size of disks and RAID group configuration for each pool.

Table 5. 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

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

Table 6. 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 7 lists the volumes provisioned for the content databases out of each Dynamic Provisioning pool.

Table 7. 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 8 lists the volumes provisioned for the search administration database out of each Dynamic Provisioning pool.

Table 8. 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 9 lists the volumes provisioned for the crawl database out of each Dynamic Provisioning pool.

Table 9. 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 10 lists the volumes provisioned for the search property database out of each Dynamic Provisioning pool.

Table 10. 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 11 lists the volumes provisioned for the `tempdb` file out of each Dynamic Provisioning pool.

Table 11. `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 12 lists the volumes provisioned for the index volumes out of each Dynamic Provisioning pool.

Table 12. 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 13 lists the volumes provisioned for the boot OS volumes out of each Hitachi Dynamic Provisioning pool.

Table 13. Boot OS Volumes

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

Create Host Groups

Before you create a host group, you must enable port security. Port security is used to isolate traffic from multiple servers and group them in common configurations. Host group members must run the same operating systems and must be connected to the same port. You can define up to 2,048 LUN paths for one host group. You can assign up to 255 host groups and up to 255 WWPNs to one Fibre Channel port.

To create a host group using Storage Navigator software, follow these steps:

1. Choose **Actions** > **Ports/Host Groups** > **Create Host Groups**.

   The **Create Host Groups** window displays.

2. Assign a name in the **Host Group Name** field.
3. From the **Host Mode** drop-down menu, choose **2C[Windows]**.
4. In the **Available Hosts** pane, highlight one or more hosts.
5. In the **Available Ports** pane, highlight one or more ports.

6. Click the **Add** button.

   The **Selected Host Groups** pane is populated.

7. Click the **Finish** button.

   The **Create Host Groups** window displays.

8. Click the **Apply** button.

Create Dynamic Provisioning Pools

To create a Dynamic Provisioning pool using Hitachi Storage Navigator software, follow these steps:

1. Choose **Actions** > **Pool** > **Create Pools**.

   The **Create Pool**s dialog box displays.

2. From the **Pool Type** drop-down menu, choose **Dynamic Provisioning**.

3. Select the **Disable** radio button for the **Multi-Tier Pool** option.

4. Choose a menu item from the **Drive Type/RPM** drop-down menu and from the **RAID Level** drop-down menu.

   These options allow you to filter the available pool volumes. For both database and log pools we utilized LDEVs from RAID 10(2D+2D) groups consisting of SAS/10K drives.

5. Click the **Select Pool VOLs** button.

   The **Select Pool VOLs** dialog box displays.

6. Highlight one or more pool volumes in the **Available Pool Volumes** pane and click the **OK** button.

   The **Create Pools** dialog box displays with **Total Selected Pool Volume** and **Total Selected Capacity** fields populated.

7. Assign a prefix for the pool name in the **Prefix** field.

   In our environment we used the following prefixes to help identify what each pool was being utilized for:

   Prefix for Database Pool: HDP-MSFT-SP-DB

   Prefix for Log Pool: HDP-MSFT-SP-LOG

8. (Optional) Assign an initial number for the pool name in the **Initial Number** field.

9. Expand the **Options** pane.

10. Assign a pool id in the **Pool ID** field.
11. Assign a subscription limit in the **Subscription Limit** field.

    This is the percentage of oversubscription for this pool that you allow in your environment.
12. Choose a value from the **User-Defined Threshold** drop-down menu and click the **Add** button.

    The **User Defined Threshold** value determines when a pool capacity alert is triggered.

    The **Selected Pools** pane is populated.
13. Click the **Finish** button.

    The **Create Pools Dialogue** box displays.
14. (Optional) Click the **Next** button to start the creation of LDEVs within the pool as well selecting the host groups to which they will be assigned.

    These operations are explained in more detail in the following sections.
15. Click the **Apply** button.

Repeat these steps to create additional pools for the environment.

Create LDEVs Within Dynamic Provisioning Pools

To create a LDEV using Storage Navigator software, follow these steps:

1. Choose **Actions** > **Logical Device** > **Create LDEVs**.

   The **Create LDEVs** window displays.
2. From the **Provisioning Type** drop-down menu, select **Dynamic Provisioning**.
3. From the **Emulation Type** drop down menu, select **OPEN-V**.
4. Choose a menu item from the **Drive Type/RPM** drop-down menu and from the **RAID Level** drop-down menu.

   For both database and log pools we utilized SAS/10K drives on a RAID 10(2D+2D) RAID level.
5. Click the **Select Pool** button.

   The **Select Pool** window displays.
6. Highlight a pool in the **Available Pools** pane and click **OK**.

   Select HDP-MSFT-SP-DB when creating the database volumes (including tempdb) as well as the index volumes for the environment. Select HDP-MSFT-SP-LOG when creating the log volumes for the environment.

   The **Create LDEVs** window displays with the **Selected Pool Name** and the **Selected Pool Capacity** fields populated.
7. Enter a capacity amount in the **LDEV Capacity** field and choose a unit of measure from the drop-down menu.

8. Enter the number of LDEVs of that size to be created in the **Number of LDEVs** field.
9. In the **LDEV Name** pane, assign a prefix in the **Prefix** field and assign an initial number in the **Initial Number** field.
10. Expand the **Options** pane.
11. Review the value in the **LDKC** field.

    Modify the LDKC value if the default of `00` is not appropriate. This is most often the case if the storage will be configured with more than one LDKC.
12. Choose a value from the **CU** drop-down menu.
13. Choose a value from the **DEV** drop-down menu.
14. (Optional) Choose a value from **Interval** drop-down menu.

    Leave this value at the default of `0` for sequential numbering of LDEVs. If you want a different numbering sequence, choose a different value.
15. Review the default values in the **Initial SSID** field, the **CLPR** field and **Processor Blade** field.

    In most situations, use the default values. Change them only if your environment requires different values.
16. Click the **Add** button.

    The **Selected LDEVs** pane is populated.
17. (Optional) Click the **Next** button to start the selection of the LDEVs that you'd like to assign to host groups.

    If you choose to execute this step, steps 18 and 19 will only be executed once you fully configure the environment. This operation is explained in more detail in the following section.
18. Click the **Finish** button.

    The **Create LDEVs** window displays.
19. Click the **Apply** button.

Map LDEVs to Host Groups

To map an LDEV using Hitachi Storage Navigator software, follow these steps:

1. Choose **Actions** > **Logical Device** > **Add LUN Paths**.

   The **Add LUN Paths** window displays.
2. In the **Available LDEVs** pane, highlight one or more LDEVs.
3. Click the **Add** button.

   The **Selected LDEVs** pane is populated.
4. Click the **Next** button.

5. In the **Available Host Groups** pane, highlight one or more host groups.
6. Click **Add**.

   The **Selected Host Groups** pane is populated.
7. Click **Next**.

   The **Add LUN Paths** window displays.
8. Click the **Finish** button.

   The **Add LUN Paths** window displays.
9. Click the **Apply** button.

Configure Servers and VMs

Install the appropriate edition of Windows Server 2008 R2 on the hosts. For the hosts being used for the SharePoint VMs, install the Datacenter edition. For the servers being used for the SQL Server instances, install the Enterprise edition.

After the operating system is installed, download and apply any available operating system patches and verify that the server's BIOS and the HBA firmware are up to date.

Install any required administrative tools such as the HBA management software that will be used. 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.

After the Hyper-V hosts are set up and up to date, follow these high-level steps to configure the Hyper-V hosts and create the VMs:

1. Set up the disk or disks that will be used to host the Hyper-V virtual machines' files.
2. Install the Hyper-V role and apply any applicable updates.
3. Create the virtual networks for the environment.
4. Create the VMs.

   Ensure that you select the appropriate location where the VHD and configuration files will be saved, and that the correct virtual networks are selected for the VM.

   Also be sure to assign the index partition volumes to the VMs (one per VM) as described in Table 12.
5. Install the appropriate OS on the VMs and run Windows update to ensure that the OS is up to date.

For more information about setting up Hyper-V and deploying virtual machines, see the Microsoft TechNet article "Hyper-V," especially the Getting Started, Planning, Installation and Configuration sections.

Run the Hyper-V Best Practices Analyzer to ensure your environment is appropriately configured. For more information, see the Hyper-V Best Practices Analyzer web site.

Install and Configure SQL Server 2008 R2

This section describes the basic requirements that must be kept in mind while installing and configuring SQL Server for SharePoint 2010 deployments.

Note that when choosing the collation type for the SQL Server instance for use with SharePoint 2010, choose the following options to match the collation type used by default for SharePoint Server:

- Latin1_General
- Case Insensitive (CS)
- Accent Sensitive (AS)
- Kana Sensitive (KS)
- Width Sensitive (WS)

The default location for the database and log files can remain unchanged at this time. After the installation is complete, the `tempdb` databases must be moved to the Virtual Storage Platform. This procedure is described in the "Configure SQL Server 2008 R2" section.

Use the SQL Server Best Practices Analyzer to analyze and identify collation selection and other best practices. For more information, see the Microsoft SQL Server 2008 R2 Best Practice Analyzer web site.

For more information, see the Microsoft TechNet article "How to: Install SQL Server 2008 R2 (Setup)."

Configure SQL Server 2008 R2

The following sections describe storage configuration requirements for SQL Server when hosting the databases, logs and `tempdb` files on the Virtual Storage Platform.

Databases and Transaction Log Files

When provisioning the storage for the SQL databases and logs, you must establish an allocation strategy using the SIZE, FILEGROWTH and MAXSIZE configuration parameters. Use the parameters illustrated in the following SQL code sample when creating the databases and transaction logs for your environment:

```
CREATE DATABASE <name of database> ON PRIMARY (
      NAME = <logical database name>,
      FILENAME = <OS location and name of database>,
      SIZE = <initial size of the database file in MB>,
      MAXSIZE = <maximum size to which the database file can grow in MB>,
      FILEGROWTH = <specifies the database growth increment in MB>)
      LOG ON (
      NAME = <logical log file name>,
      FILENAME = <OS location and name of log file>,
      SIZE = <size for log file in MB>,
      MAXSIZE = <maximum size to which the log file can grow in MB>,
      FILEGROWTH <specifies the log growth increment in MB>)
```

For the implementation described in this white paper, the following parameters were used (placeholder variables are replaced with lab-specific parameters and are shown in red):

```
CREATE DATABASE mydb ON PRIMARY (
      NAME = mydb,
      FILENAME = 'E:\mountpoint\data1\mydb1.mdf',
      SIZE = 420MB,
      MAXSIZE = 42000,
      FILEGROWTH = 42)
      LOG ON (
      NAME = mydblog1,
      FILENAME = 'L:\mountpoint\log\mydblog1.ldf',
      SIZE = 420MB,
      MAXSIZE = 10000,
      FILEGROWTH = 42)
```

Each file is created with an initial SIZE and auto extends by FILEGROWTH when the currently allocated space runs out, but stops if it reaches MAXSIZE. For the database, the values of SIZE, FILEGROWTH and MAXSIZE are determined by expected growth. Make sure that the value for FILEGROWTH is a multiple of a page, 42MB, which is ideal when using Hitachi Dynamic Provisioning software. MAXSIZE is more straightforward, and is used to make sure that your application does not exceed capacity, fail or both. Set SIZE to a value that creates an initial allocation of data that is at least equal to the amount of data you plan to import into the database or that you expect the database or log files to be initially. Pay special attention to the SIZE and FILEGROWTH values given to log files, as small sizes might affect system performance and the log files may grow to a large size due to many small increments. This can slow database startup as well as log backup and restore operations for a given SQL Server instance. Microsoft recommends that you assign log files a SIZE value close to the final size required, and also have a relatively large FILEGROWTH value (use a multiple of 42MB for implementations that use Hitachi Dynamic Provisioning software).

Use the instant file initialization feature of SQL Server to allow faster and optimized data file creation and growth and fast execution of database or filegroup restore operations. Instant file initialization reclaims used disk space without filling that space with zeros. Instead, disk content is overwritten as new data is written to the files, which makes it an ideal SQL Server feature to work in conjunction with Hitachi Dynamic Provisioning software. For more information, see the Data File Initialization page of SQL Server 2008 Books Online.

tempdb Files

The default location for tempdb files in SQL Server is on the local server `C:` drive. You can modify the existing `tempdb` file location to point to the LUN provisioned on the Virtual Storage Platform. To do this, execute the following SQL statements:

```
USE tempdb
GO

ALTER DATABASE tempdb
MODIFY FILE (NAME='tempdev', FILENAME= <OS location and name of tempdb>,
SIZE = <size of file in KB or MB>,
FILEGROWTH = <specifies the tempdb growth increment in MB>)
```

For the implementation described in this white paper, the following parameters were used (placeholder variables are replaced with lab-specific parameters and are shown in red):

```
USE tempdb
GO

ALTER DATABASE tempdb
MODIFY FILE (NAME='tempdev', FILENAME= 'T:\SQL\DATA\datatempdb.mdf',
SIZE=8192KB, MAXSIZE=120GB FILEGROWTH=42MB)


ALTER DATABASE tempdb
MODIFY FILE (NAME='templog', FILENAME= 'U:\SQL\DATA\datatemplog.ldf',
SIZE=8192KB, MAXSIZE=120GB FILEGROWTH=42MB)
```

Note that the files are now located on a LUN that is assigned to the SQL Server for the `tempdb`.

Create additional tempdb files for each CPU core that the server has. The Hitachi Data Systems test environment had a total of four CPU cores on the test server. For the implementation described in this white paper, the following parameters were used (placeholder variables are replaced with lab-specific parameters and are shown in red):

```
USE tempdb
GO

ALTER DATABASE tempdb
ADD FILE (NAME='tempdev1', FILENAME= 'U:\SQL\DATA\datatempdb1.mdf',
SIZE=8192KB, MAXSIZE=120GB FILEGROWTH=42MB)


ALTER DATABASE tempdb
ADD FILE (NAME='tempdev2', FILENAME= 'T:\SQL\DATA\datatempdb2.mdf',
SIZE=8192KB, MAXSIZE=120GB FILEGROWTH=42MB)


ALTER DATABASE tempdb
ADD FILE (NAME='tempdev3', FILENAME= 'U:\SQL\DATA\datatempdb2.mdf',
SIZE=8192KB, MAXSIZE=120GB FILEGROWTH=42MB)
```

The procedure described in this section can also be used to relocate the SharePoint Server database and log files in case they are accidentally created in the wrong location.

For more information about deploying using DBA-created databases, see the Microsoft TechNet article "Deploy by using DBA-created databases (SharePoint 2010)." For more information about configuring SQL Server for SharePoint 2010 deployments, see the Microsoft TechNet article "SQL Server and storage (SharePoint 2010)."

Install and Configure SharePoint Server 2010

The deployment and configuration for your SharePoint environment largely depends on the features you plan to install in your SharePoint environment and on how you plan to deploy the environment. This implementation guide focuses on the key components that an enterprise-class SharePoint environment must have to support a collaboration environment.

This implementation guide concentrates on the tasks related to the storage configuration for SharePoint 2010 deployments. For more information about how to install and configure each of the components for your SharePoint deployment, see the Microsoft TechNet article "Deployment for SharePoint Server 2010."

After the SharePoint configuration is finalized, configure the network load balance setup for your environment. For more information, see the Microsoft TechNet article "Network Load Balancing Deployment Guide."

Replicate Active Databases to Mirror SQL Server

Additional Dynamic Provisioning pools and LDEVs must be created and configured to replicate the active databases to the second SQL Server. After the LDEVs are created, follow the standard SQL Server database mirroring steps to replicate the active databases being used for the SharePoint Server environment.

For more information about replicating the active databases being used for your SharePoint 2010 environment, see the Microsoft TechNet article, "Database Mirroring Deployment."

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