HITACHI
Inspire the Next

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

Implementation Guide

By Mark Crowley

July 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

Deploying Microsoft® SharePoint® Server 2010 on the Hitachi Virtual Storage Platform using Hitachi Dynamic Tiering

Implementation Guide

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

Hitachi simplifies this task with the Hitachi Virtual Storage Platform, an enterprise-class storage system. The Virtual Storage Platform features Hitachi Dynamic Provisioning and Hitachi Dynamic Tiering, which supports thin provisioning, wide striping, and automatically tiered storage functionality.

In a SharePoint 2010 environment, systems administrators constantly deal with performance and sizing challenges. These relate to the multitude of documents and indexed data that is typically referred to as “SharePoint sprawl.” The Hitachi Virtual Storage Platform addresses these challenges with one storage platform.

This implementation guide gives database administrators the information to successfully plan and deploy a Microsoft SharePoint environment. It describes the interaction between the many components of Microsoft SharePoint Server and Microsoft SQL Server. It also describes how to use the Hitachi Virtual Storage Platform to create a robust, space efficient, and highly available SharePoint environment.

This implementation guide focuses on planning and deploying a 200,000-user Microsoft SharePoint 2010 farm. The intended use is by IT personnel responsible for planning and deploying SharePoint infrastructure and associated storage. It assumes some familiarity with SAN storage, Hitachi Compute Blade 2000, Microsoft Windows 2008 R2, Microsoft Hyper-V, and Microsoft SharePoint 2010.

Solution Overview

This is the storage configuration for Microsoft SharePoint 2010 data on a Hitachi Virtual Storage Platform.

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

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

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

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

Figure 1 shows the infrastructure used to host the Microsoft SharePoint 2010 farm.

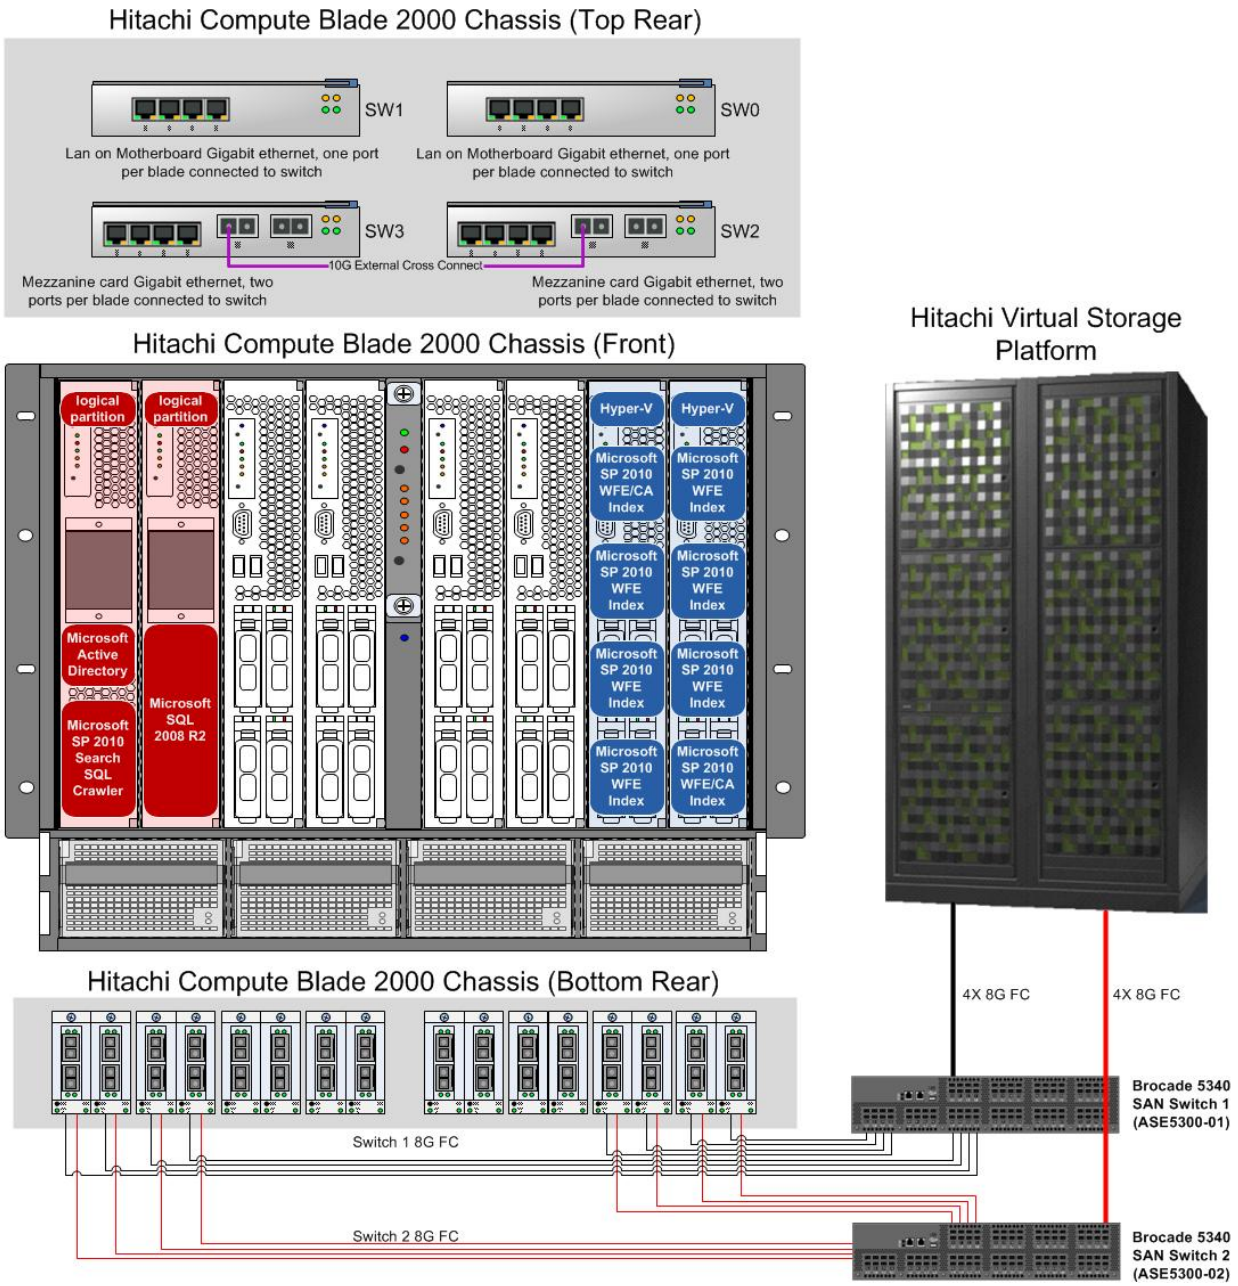
Hitachi Compute Blade 2000 Chassis (Top Rear)
SW1
SW0
Lan on Motherboard Gigabit ethernet, one port per blade connected to switch
Lan on Motherboard Gigabit ethernet, one port per blade connected to switch
SW3
SW2
10G External Cross Connect
Mezzanine card Gigabit ethernet, two ports per blade connected to switch
Mezzanine card Gigabit ethernet, two ports per blade connected to switch
Hitachi Compute Blade 2000 Chassis (Front)
logical partition
logical partition
Hyper-V
Hyper-V
Microsoft SP 2010 WFE/CA Index
Microsoft SP 2010 WFE Index
Microsoft SP 2010 WFE Index
Microsoft SP 2010 WFE Index
Microsoft Active Directory
Microsoft SQL 2008 R2
Microsoft SP 2010 Search SQL Crawler
Microsoft SP 2010 WFE Index
Microsoft SP 2010 WFE Index
Microsoft SP 2010 WFE Index
Microsoft SP 2010 WFE/CA Index
Hitachi Virtual Storage Platform
Hitachi Compute Blade 2000 Chassis (Bottom Rear)
4X 8G FC
4X 8G FC
Brocade 5340 SAN Switch 1 (ASE5300-01)
Switch 1 8G FC
Switch 2 8G FC
Brocade 5340 SAN Switch 2 (ASE5300-02)

Figure 1

Key Solution Components

The following are the key components used for the solution.

Hitachi Virtual Storage Platform

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

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

The switch matrix architecture of the Virtual Storage Platform 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 without any disruption to connected host servers. All these resources are coupled tightly through a global cache that creates a common pool of storage resources. These resources can include external storage that is connected through front-end director initiator ports.

For more information, see the Hitachi Virtual Storage Platform on the Hitachi Data Systems website.

Hitachi Dynamic Provisioning

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

Using Hitachi Dynamic Provisioning is similar to using a host-based logical volume manager (LVM), but without incurring host processing overhead. It provides one or more wide-striping pools across many RAID groups within a Hitachi Virtual Storage Platform. Each pool has one or more dynamic provisioning virtual volumes (DP-VOLs) of a user-specified logical size of up to 60TB created against it (with no initial physical space allocated).

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

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

For more information, see the Hitachi Dynamic Provisioning datasheet and Hitachi Dynamic Provisioning on the Hitachi Data Systems website.

Hitachi Dynamic Tiering

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

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

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

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

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

For more information, see Hitachi Dynamic Tiering on the Hitachi Data Systems website.

Hitachi Compute Blade 2000

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

- A balanced system architecture that eliminates bottlenecks in performance and throughput
- Embedded logical partition virtualization
- Configuration flexibility
- Eco-friendly power-saving capabilities
- Fast server failure recovery using a N+1 cold standby design that allows replacing failed servers within minutes

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

Hitachi Dynamic Link Manager Advanced

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

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

Configuration of Hitachi Dynamic Link Manager Advanced, used for SAN multipathing, used its round-robin load balancing policy. This policy selects a path by rotating through all available paths. Balancing the load across all available paths optimizes IOPS and response time.

For more information, see Hitachi Dynamic Link Manager on the Hitachi Data Systems website.

Microsoft SharePoint Server 2010

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

Microsoft SQL Server 2008 R2

For more information about Microsoft SQL Server 2008 R2, see the Microsoft SQL Server web site.

Tested Solution Components

The following sections give the details the hardware and software components to deploy to support the solution described in this guide.

Hardware Components

Table 1 lists the configuration of each physical server.

Table 1. Physical Server Configuration

Server Hardware	*Firmware Level*	*Qty*	*Operating System*	*Purpose*
Hitachi Compute Blade 2000 E57A1 • 2 Intel Xeon X7560 at 2.26GHz • 12 Cores, 96GB RAM • 2 Hitachi PCIe 8Gb/sec dual port HBA • 1 Hitachi mezzanine 4 port Gigabit Ethernet Card	EFI BIOS v4.6.3	2	Microsoft Windows Server 2008 R2 Enterprise Edition	Active Directory logical partition Microsoft SharePoint 2010 search, crawl, and SQL server logical partition Microsoft SQL Server logical partition See Table 2 for configuration details

Server Hardware	*Firmware Level*	*Qty*	*Operating System*	*Purpose*
Hitachi Compute Blade 2000 E55A2 • 2 Intel Xeon X5670 at 2.93GHz • 72GB RAM • 2 Emulex LPe 12002-M8 HBA • 1 Hitachi mezzanine 4 Port Gigabit Ethernet card	EFI BIOS v4.6.3	2	Microsoft Windows Server 2008 R2 Datacenter Edition	Microsoft Hyper-V host servers See Table 3 for configuration details

Table 2 has the configuration of the logical partitions.

Table 2. Logical Partition Configuration

Machine name *Blade* *Logical Partition*	*Hardware assignment*	*Operating System* *Software Installed*
DC1.mstest.local • Blade 0 • Logical Partition 3	4 processors 6GB RAM 2 shared gigabit Ethernet ports 2 shared Fibre Channel ports	Microsoft Windows 2008 R2, Enterprise Edition, Service Pack 1 Hitachi Dynamic Link Manager Active Directory role (2008 R2 forest level), DNS role
MSFT-SP-WS-SC • Blade 0 • Logical Partition 4	4 processors 32GB RAM 2 shared gigabit Ethernet ports 2 dedicated gigabit Ethernet ports 4 shared Fibre Channel ports	Microsoft Windows Server 2008 R2, Enterprise Edition, Service Pack 1 Microsoft SharePoint 2010 Microsoft SQL Server 2008 R2 Hitachi Dynamic Link Manager
MSFT-SP-SQL-1 • Blade 1 • Logical Partition 1	12 Processors 96GB RAM 4 dedicated gigabit Ethernet ports 4 shared Fibre Channel ports	Microsoft Windows Server 2008 R2, Enterprise Edition, Service Pack 1 Microsoft SQL Server 2008 R2, Enterprise Edition Hitachi Dynamic Link Manager

Table 3 lists the Microsoft Hyper-V hosts and virtual machines used for the SharePoint Farm.

Table 3. Microsoft Hyper-V Server Configuration

Server	*VMs*	*Drives*	*Details*	*Name*	*Role*
Blade6-HV	4	Two 150GB LUNs per Virtual Machine • 1 for OS • 1 for local index	Microsoft Windows 2008 R2 Enterprise Edition, 4 VCPUs, 8GB RAM	SP-WS01 SP-WS02 SP-WS03 SP-WS04	Web front end Search index Mirror
Blade7-HV	4	Two 150 GB LUNs per Virtual Machine • 1 for OS • 1 for local index	Microsoft Windows 2008 R2 Enterprise Edition, 4 VCPUs, 8GB RAM	SP-WS05 SP-WS06 SP-WS07 SP-WS08	Web front end Search index Mirror

Figure 2 shows the internal networking configuration for the Microsoft Hyper-V hosts.

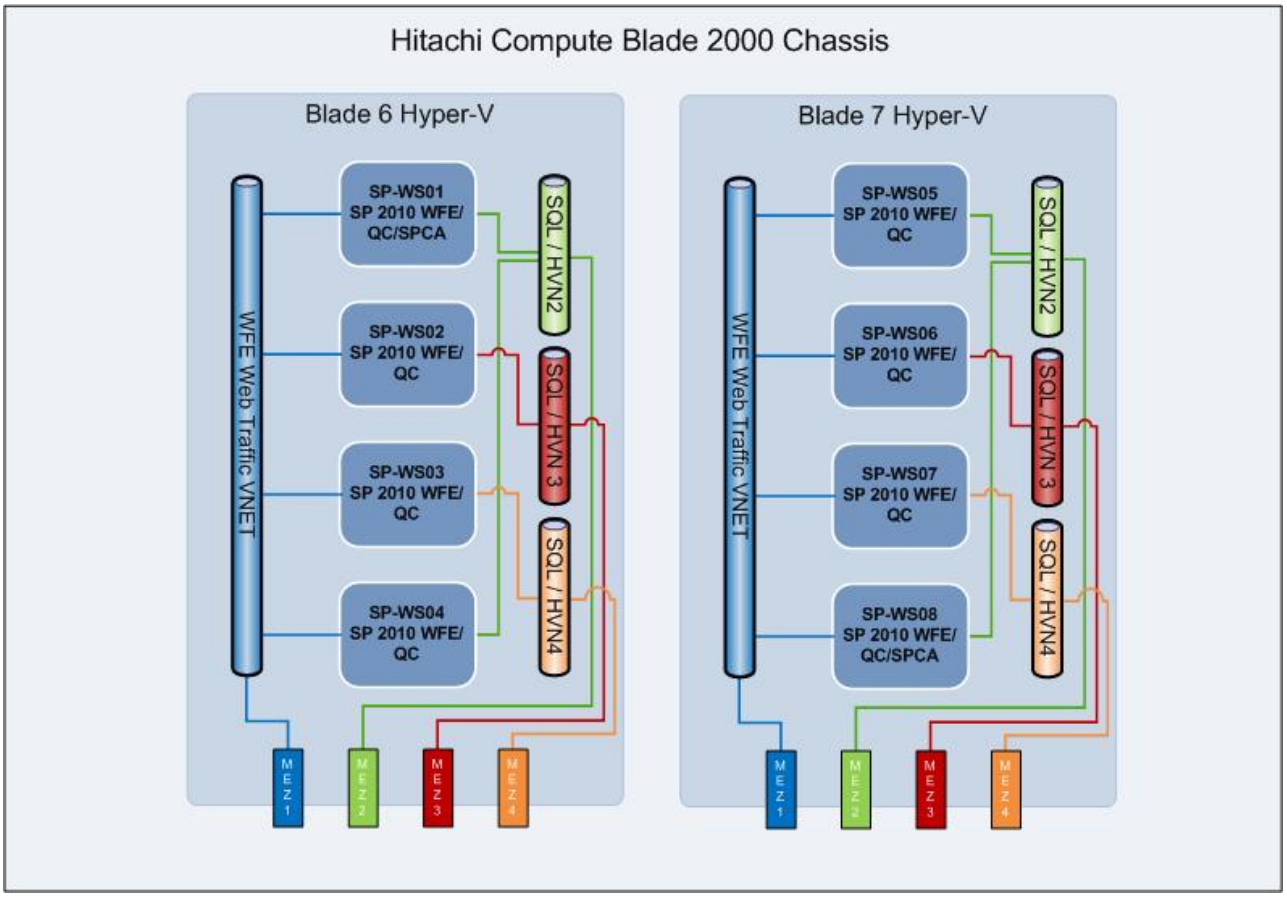

Figure 2

In Microsoft Hyper-V, a virtual switch is configured for each physical network port that is used on each server. The web front-end traffic was isolated to a single virtual switch. The SQL backend traffic was isolated to the remaining three virtual switches.

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 for all Fibre Channel connections through these switches. The various paths from the servers to the storage system are managed using Hitachi Dynamic Link Manager.

The logical partitions use virtual HBA ports, using NPIV. These HBAs are connected through the hypervisor to the physical ports. The servers containing the logical partitions each have two dual port HBAs installed in the PCIe slots. The initiators, both physical and NPIV were each zoned to individual ports on the Hitachi Virtual Storage Platform. Both SQL servers use target ports 3A, 4A, 7A, and 8A on the storage array.

Each Microsoft Hyper-V host has two dual port HBAs. However, one port from each HBA is zoned in the switch to target ports 3B and 7B on the storage array.

Eight separate LDEVs are mapped to each Microsoft Hyper-V Host. Four LDEVs are pass-through disks for the drive C: volumes on the virtual machines and the other four LDEVs are pass-through disks for the local index volumes that the virtual machines used.

Hitachi Dynamic Link Manager uses the Microsoft MPIO framework. This further optimizes the multipathing strategy for best performance.

The “extended round-robin” load balance policy is selected for the paths on all systems. This policy provides the best path utilization coupled with automatic optimization of sequential loads. The round-robin load-balancing algorithm automatically selects a path by rotating through all available paths to balance the load across all available paths and optimize IOPS and response time.

Figure 3 diagrams Microsoft SharePoint SQL, Microsoft Hyper-V storage infrastructure, and Fiber Channel connectivity. The term “PTD” in Figure 3 refers to the use of *pass-through disks* in Hyper-V for the guest virtual machines.

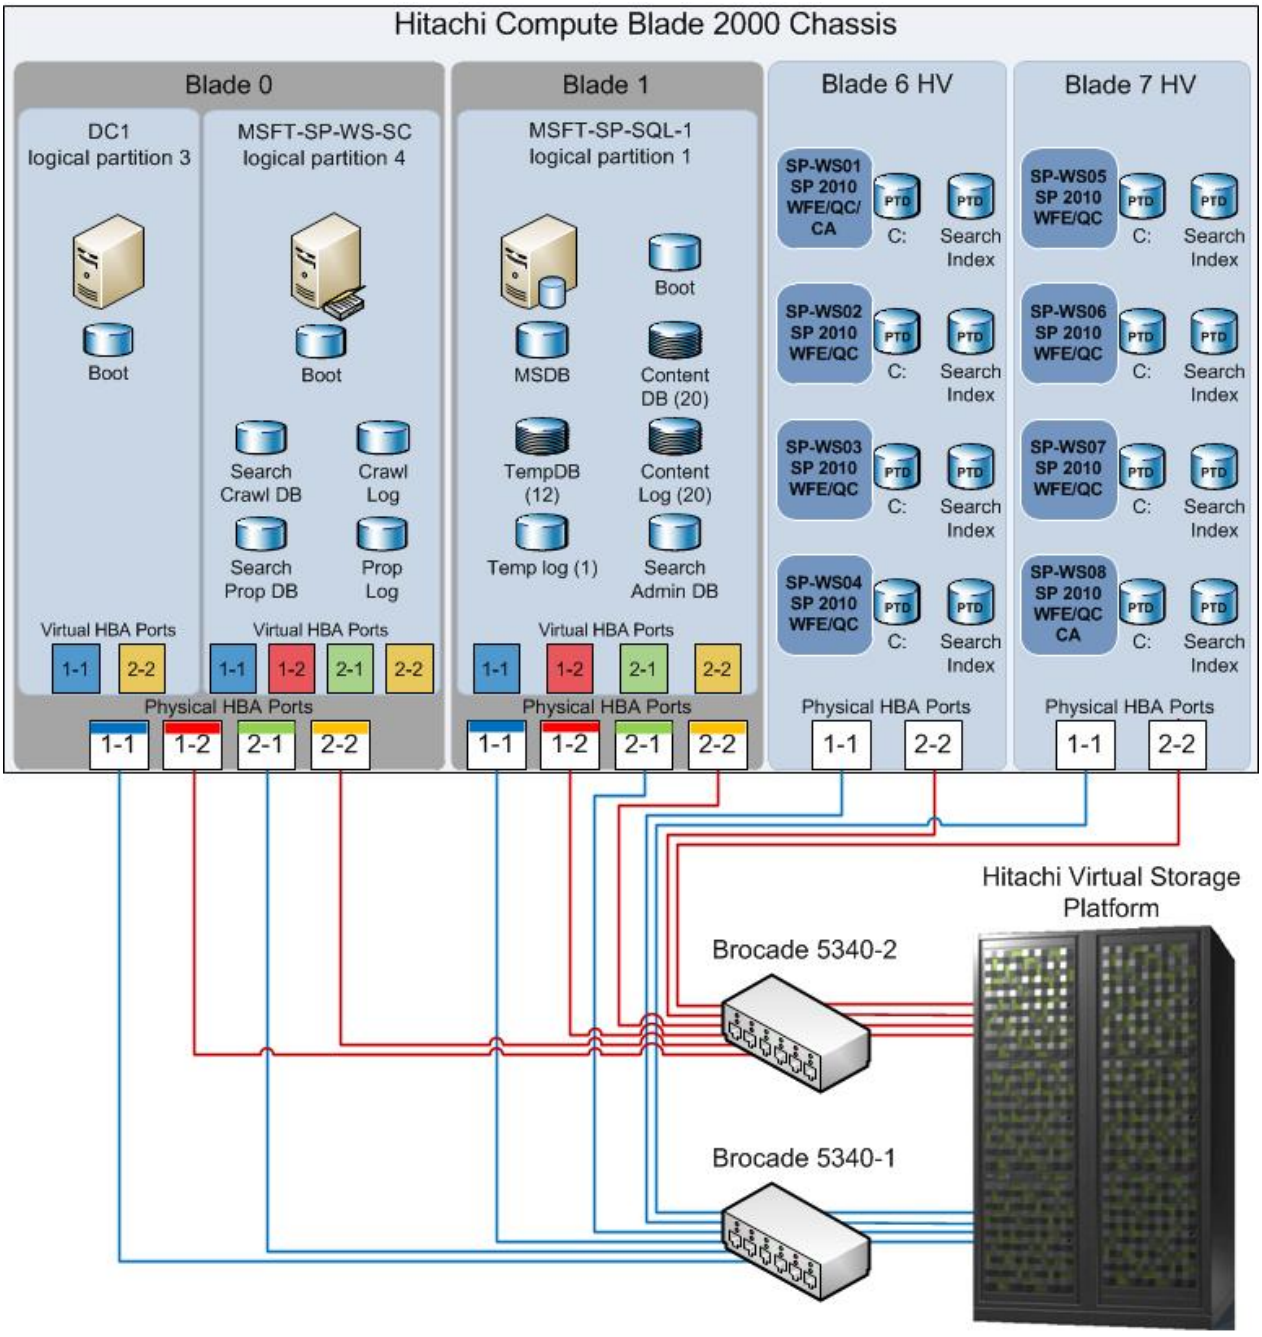

Figure 3

Software Components

Table lists the detailed information about the software components used.

Table 4. Software Components

Software	*Version*
Hitachi Storage Navigator	Dependent on microcode version
Hitachi Dynamic Provisioning	Dependent on microcode version
Hitachi Dynamic Tiering	Dependent on microcode version
Hitachi Dynamic Link Manager	6.5.0-00
Microsoft Windows Server 2008 R2, Datacenter edition	6.1.7601 (SP1)
Microsoft Windows Server 2008 R2, Enterprise edition	6.1.7601 (SP1)
Microsoft SQL Server 2008 R2	Enterprise edition
Microsoft SharePoint 2010	Enterprise edition

Solution Implementation

To deploy this Microsoft SharePoint Server 2010 solution, follow these steps:

1. Configure fabric switch zones.
2. Configure the Hitachi Virtual Storage Platform.
3. Configure the servers.
4. Enable the Hyper-V role on the Virtualization Hosts.
5. Deploy the virtual machines.
6. Install and configure Microsoft SQL Server 2008 R2.
7. Install and configure Microsoft SharePoint Server 2010.

These are the general tasks necessary for a successful deployment. Your checklist might vary from these steps based on your environment and needs. The following sections give the details about each of the above steps.

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

- Hitachi Storage Navigator online help
- *Hitachi Virtual Storage Platform Provisioning Guide for Open Systems*
- 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 following 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 lists the zoning details for the SAN.

Table 5. SAN Switch Architecture

Server	*HBA Ports*	*Switch Zone*	*Storage Port*	*Switch*
Blade6-HV	HBA1-1	BS2K_01_B6_HBA1_1_ASE45_36_3B	3B	5300-01
Blade6-HV	HBA2-2	BS2K_01_B6_HBA2_2_ASE45_36_7B	7B	5300-02
Blade7-HV	HBA1-1	BS2K_01_B7_HBA1_1_ASE45_36_3B	3B	5300-01
Blade7-HV	HBA2-2	BS2K_01_B7_HBA2_2_ASE45_36_7B	7B	5300-02
MSFT-SP-SQL-1	HBA1-1	BS2K_01_B1_L1_HBA1_1_ASE45_36_3A	3A	5300-01
MSFT-SP-SQL-1	HBA1-2	BS2K_01_B1_L1_HBA2_1_ASE45_36_4A	4A	5300-01
MSFT-SP-SQL-1	HBA2-1	BS2K_01_B1_L1_HBA1_2_ASE45_36_7A	7A	5300-02
MSFT-SP-SQL-1	HBA2-2	BS2K_01_B1_L1_HBA2_2_ASE45_36_8A	8A	5300-02
MSFT-SP-WS-SC	HBA1-1	BS2K_01_B0_L4_HBA1_1_ASE45_36_3A	3A	5300-01
MSFT-SP-WS-SC	HBA1-2	BS2K_01_B0_L4_HBA1_2_ASE45_36_4A	4A	5300-01
MSFT-SP-WS-SC	HBA2-1	BS2K_01_B0_L4_HBA2_1_ASE45_36_7A	7A	5300-02
MSFT-SP-WS-SC	HBA 2-2	BS2K_01_B0_L4_HBA2_2_ASE45_36_8A	8A	5300-02
SANBOOT DC1	HBA1-1	BS2K_01_B0_L3_HBA1_1_ASE45_36_4B	4B	5300-01
DC1	HBA2-2	BS2K_01_B0_L3LHBA2_2_ASE45_36_7B	7B	5300-02
SANBOOT MSFT-SP-SQL-1	HBA1-1	BS2K_01_B1_L1_HBA1_1_ASE45_36_4B	4B	5300-01
SANBOOT MSFT-SP-WS-SC	HBA1-1	BS2K_01_B0_L4_HBA1_1_ASE45_36_4B	4B	5300-01

Configure the Hitachi Virtual Storage Platform

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

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

Table 6. Dynamic Provisioning Pools

Pool Name (ID)	*RAID Group Configuration*	*Drive Type*	*Number of RAID Groups*	*Pool Capacity (TB)*
HDT-MSFT-SP-DB(MT) (30)	RAID-5 (7D+1P)	300GB, 10K RPM, SAS	1	1.88
		2TB, 7.2K RPM, SATA	2	5.14
HDT-MSFT-SP-LOG (31)		300GB, 10K RPM, SAS	1	1.88
HDT-MSFT-SP-VMDATA (35)		300GB, 10K RPM, SAS	1	1.88

Table lists the volumes provisioned for the central administration database out of each dynamic provisioning pool.

Table 7. Content Databases Volumes Presented to MSFT-SP-SQL-1

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

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

Table 8. Search Administration Database Volumes Presented to MSFT-SP-SQL-1

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

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

Table 9. Crawl Database Volumes Presented to MSFT-SP-WS-SC

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

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

Table 10. Search Property Database Volumes Presented to MSFT-SP-WS-SC

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

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

Table 4. tempDB Volumes Presented to MSFT-SP-SQL-1

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

Table lists the volumes provisioned for the Microsoft Hyper-V virtual machines and index volumes.

Table 5. Virtual Machine Boot and Index Volumes

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

Create Host Groups

Enable port security before you create a host group.

Port security does the following:

- Isolate traffic from multiple servers
- Groups servers in common configurations.

Host group members must use the same operating system and must connect to the same port. You can do the following:

- Define up to 2,048 LUN paths for one host group.
- Assign up to 255 host groups and up to 255 WWPNs to one Fibre Channel port.

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

1. From the **Actions** menu, point to **Ports/Host Groups** and click **Create Host Groups**.

   The **Create Host Groups** window opens.

2. Type a name for the group in **Host Group Name**.
3. From the **Host Mode** list, click **2C[Windows]**.
4. In the **Available Hosts** area, select one or more hosts.

5. In the **Available Ports** pane, select one or more ports.
6. Click **Add**.

   The **Selected Host Groups** area is populated with the selected port.
7. Click **Finish**.

   The **Create Host Groups** window opens.
8. Click **Apply**.

Create Dynamic Provisioning Pools

To create a dynamic provisioning pool using Hitachi Storage Navigator, select the pool volumes and then create the dynamic provisioning pool. Complete both of the following.

To select the pool volumes for the dynamic provisioning pool:

1. From the **Actions** menu, point to **Pool** and then click **Create Pools**.

   The **Create Pool**s dialog box opens.
2. From the **Pool Type** list, click **Dynamic Provisioning**.
3. Click **Enable** for the **Multi-Tier Pool** option.
4. To filter the available pool volumes, click an item from the **Drive Type/RPM** list and from the **RAID Level** list.

   The database pool in this solution used the following:

   - Two RAID 5 (7D+1P) groups of 2TB SATA drives. When provisioning these RAID 5 groups, the system divided them into five units, each with 20% of the total RAID group capacity.
   - One RAID 5 (7D+1P) group of 300GB SAS drives.

   The SQL, log, and VM pools in this solution each used an LDEV from RAID 5(7D+1P) groups with SAS drives. See Table .
5. Click **Select Pool VOLs**.

   The **Select Pool VOLs** dialog box opens.
6. Select one or more pool volumes in the **Available Pool Volumes** area and click **OK**.

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

To create the dynamic provisioning pool, follow these steps.

1. Type a prefix for the pool name in the **Prefix** field.

   The Hitachi Data Systems lab environment used these prefixes to identify the pools shown in Table :

   - Prefix for database pool: **HDP-MSFT-SP-DB(MT)**
   - Prefix for log pool: **HDP-MSFT-SP-LOG**
   - Prefix for virtual machine storage: **HDP-MSFT-SP-VMDATA**
2. (Optional) Type the initial number for the first pool name in **Initial Number**.
3. Expand the **Options** area.
4. Type a pool ID in **Pool ID**.
5. Assign a subscription limit in **Subscription Limit**.

   This sets the percentage of oversubscription allowed for this pool in your environment.
6. Click a value from the **User-Defined Threshold** list and then click **Add**.

   The **User Defined Threshold** value determines when to trigger a pool capacity alert.

   The **Selected Pools** pane is populated with the pool(s) that are being created.
7. Click **Finish**.

   The **Create Pools Dialogue** box opens.
8. (Optional) Click the **Next** button to start the following:

   - The creation of LDEVs within the pool. See “Create LDEVs Within the Pools.”
   - The selection of the host groups to assign the LDEVs. See “Map LDEVs to Host Groups.”
9. Click **Apply**.

Repeat the steps for selecting and creating the dynamic provisioning pools to create enough pools for the environment.

Create LDEVs Within the Pools

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

1. In Storage Navigator on the **Explorer** column, select **Pools.** Click the name of the pool in which to create LDEVs.
2. Click **Create LDEVs**.
3. Type a size for the LDEV(s) in **LDEV Capacity:** and the number of LDEVs in **Number of LDEVs:**.
4. Type a prefix for the LDEV or LDEVs that you are creating in **Prefix.** This would be something like *HDT-MSFT-SP-CDB*.

5. Type the starting number in **Initial Number** if creating more than one LDEV (for example, **01**). This number increases by 1 for each LDEV created.
6. Expand the **Options** section under the **LDEV Name:** section.
7. In the **CU** list, click the number of the pool that you are using (for example, **30**). Use the pool ID to identify the source pool for the LDEVs more easily.
8. In the **DEV** list, click the starting number of the LDEV that you are creating. You cannot create more than one LDEV with the same **CU:DEV** combination.
9. Click **Add->**. The created LDEV or LDEVs populate the **Selected LDEVs** column.
10. Click **Finish** and then click **Apply**.

Storage Navigator will generate an execution task and start on it. You can select the **Tasks** item in the **Explorer** column to see about task status.

Map LDEVs to Host Groups

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

1. From the **Actions** menu, point to **Logical Device** and then click **Add LUN Paths**.

   The **Add LUN Paths** window opens.
2. In the **Available LDEVs** area, select one or more LDEVs.
3. Click **Add**.

   The **Selected LDEVs** area is populated with the selected LDEVs.
4. Click **Next**.
5. In the **Available Host Groups** area, select one or more host groups.
6. Click **Add**.

   The **Selected Host Groups** pane is populated with the selected host groups.
7. Click **Next**.

   The **Add LUN Paths** window opens.
8. Click **Finish**.

   The **Add LUN Paths** window opens.
9. Click **Apply**.

Configure the servers

Install the appropriate edition of Microsoft Windows Server 2008 R2 on the hosts.

- **Datacenter edition** — For the Microsoft Hyper-V hosts.
- **Enterprise edition** — For all other servers and virtual machines.

After installing the operating system, do the following:

- Download and apply all available operating system patches.
- Verify that the server's BIOS and the HBA firmware are up to date.

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

When choosing Fibre Channel card drivers, verify that you are using the current recommended drivers for the Hitachi Virtual Storage Platform. Major vendors have web pages for you to download the current Hitachi storage systems drivers. For a list of currently supported Fibre Channel cards and drivers, see "Interoperability" on the Hitachi Data Systems web site.

Enable the Hyper-V Role on the Virtual Hosts

Follow these steps to enable the Hyper-V role on the virtual hosts:

1. Open **Server Manager**.
2. Click **Roles**. To the right, click the **Add Roles** wizard.

   An **Add Roles Wizard** window opens.
3. Click **Next**.

   A list of available roles will show up.
4. Select the **Hyper-V** check box, and then click **Next**.
5. Answer the questions for the system's configuration options, such as network. When finished, click to confirm those answers.

   Microsoft Windows then reboots twice and installs the hypervisor. When you are able to log on to the system, Microsoft Hyper-V is installed.

Refer to the Microsoft *TechNet* article "Hyper-V" for more details on the deployment and operation of Hyper-V on Microsoft Windows Server 2008 R2.

Deploy the Virtual Machines

Administrators have various options for the deployment of the virtual machines within a Microsoft Hyper-V environment. Use your method to deploy the eight virtual machines on the two Hyper-V enabled hosts.

Refer to Figure 2 for a suggested networking configuration for the Hyper-V hosts. Microsoft network load balancing was used to load balance the network adapters connected to the **WFE Web Traffic VNET** in this example.

For more information about setting up Microsoft Hyper-V to deploy virtual machines, see "Hyper-V" on Microsoft *TechNet,* especially Getting Started, Planning, Installation, and Configuration sections.

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

Install and Configure Microsoft SQL Server 2008 R2

Keep these requirements in mind while installing and configuring Microsoft SQL Server 2008 R2 for Microsoft SharePoint 2010 deployments.

When choosing the collation type for the SQL server instance for use with Microsoft SharePoint 2010, these options match the default collation type used for Microsoft SharePoint:

- 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 completing the installation, move the `tempdb` databases to the Hitachi Virtual Storage Platform, as described in "Configure SQL Server 2008 R2."

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

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 Microsoft SQL Server when hosting the databases, logs and `tempdb` files on the Hitachi 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. For the database, the values of SIZE, FILEGROWTH and MAXSIZE are determined by expected growth.

- Each file is created with a default initial SIZE. At a minimum, set this value to be large enough to handle either of the following:
  - The initial allocation of data that you plan to import into the database
  - The initial size that you expect for the database or log files.
- In this solution, the minimum file size for all of the databases was altered to be 420MB, and the logs to 42MB.
- The file auto extends in size by FILEGROWTH when the currently allocated space for the file runs out. Use a value for FILEGROWTH that is a multiple of 42MB. Use a multiple of this size because, when using Hitachi Dynamic Provisioning and Hitachi Dynamic Tiering, the size of a page is 42MB.
- The file growth stops if it reaches MAXSIZE. This is used to make sure that your application does not exceed capacity of the LUN.

Pay special attention to the SIZE and FILEGROWTH values given to log files. Small sizes might affect system performance. The log files may grow to a large size resulting from many small increments. This can slow database startup as well as log backup and restore operations for a given SQL server instance. Microsoft recommends the following:

- Assign log files a SIZE value close to the *final* size required (not initial)
- Have a relatively large FILEGROWTH value. Remember to use a multiple of 42MB for implementations that use Hitachi Dynamic Provisioning.

Use the below TSQL code samples to create databases on your SQL server.

To do this, execute the following SQL statements in the query window opened by the **New Query** option in SQL server 2008 R2 Management Studio)

```
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 kb/mb/gb/tb>,
      MAXSIZE = <the size limit of the database file in kb/mb/gb/tb>,
      FILEGROWTH = <the database growth increment in kb/mb/gb/tb>)
      LOG ON (
      NAME = <logical log file name>,
      FILENAME = <OS location and name of log file>,
      SIZE = <initial size for log file in kb/mb/gb/tb>,
      MAXSIZE = <the size limit of the log file in kb/mb/gb/tb>,
      FILEGROWTH <specifies the log growth increment in kb/mb/gb/tb>)
```

For the implementation described in this installation guide, the following parameters were used, replacing the placeholder variables with lab-specific parameters shown in red:

```
CREATE DATABASE COL01 ON PRIMARY (
      NAME = COL01,
      FILENAME = 'C:\SQLDBMNT\COL01DB\DB\COL01.mdf',
      SIZE = 420mb,
      MAXSIZE = 200000mb,
      FILEGROWTH = 42mb)
      LOG ON (
      NAME = COLO1_LOG,
      FILENAME = 'C:\SQLLOGMNT\COL01LOG\COL01.ldf',
      SIZE = 420mb,
      MAXSIZE = 20000mb,
      FILEGROWTH = 42mb)
```

Use the instant file initialization feature of Microsoft SQL Server to allow the following:

- Faster and optimized data file creation and growth
- Faster execution of database or file group restore operations

Instant file initialization reclaims used disk space without filling that reclaimed space with zeros. Instead, disk content is overwritten as new data is written to the files. This makes it an ideal SQL server feature to work in conjunction with Hitachi Dynamic Provisioning. 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 drive C of the server. Use the ALTER DATABASE TSQL command to modify the `tempDB` file location to the LUN provisioned for that purpose on the Hitachi Virtual Storage Platform. To do this, execute the following SQL statements in the query window opened by the **New Query** option in SQL server 2008 R2 Management Studio)

```
USE tempdb
GO

ALTER DATABASE tempdb
MODIFY FILE (NAME='tempdev', FILENAME= <OS location and name of tempdb>,
SIZE = <initial size of the database file in kb/mb/gb/tb>,
MAXSIZE = <the size limit of the database file in kb/mb/gb/tb>,
FILEGROWTH = <the database growth increment in kb/mb/gb/tb>)

ALTER DATABASE tempdb
MODIFY FILE (NAME='templog', FILENAME= <OS location and name of tempdb>,
SIZE = <initial size of the log file in kb/mb/gb/tb>,
MAXSIZE = <the size limit of the log file in kb/mb/gb/tb>,
FILEGROWTH = <the log growth increment in kb/mb/gb/tb>)
```

For the implementation described in this installation guide, the following parameters were used, replacing the placeholder variables with lab-specific parameters shown in red:

```
USE tempdb
GO

ALTER DATABASE tempdb
MODIFY FILE (NAME='tempdev', FILENAME= 'C:\SQLDBMNT\TempDB\tempdb.mdf',
SIZE=420MB, MAXSIZE=120GB, FILEGROWTH=42MB)


ALTER DATABASE tempdb
MODIFY FILE (NAME='templog', FILENAME= 'C:\SQLLOGMNT\TempLog\templog.ldf',
SIZE=42MB, MAXSIZE=20GB, FILEGROWTH=42MB)
```

Note — Set the file sizes and growth increment to a multiple of 42MB to align sizes used in the databases with the page size of the Hitachi Virtual Storage Platform.

This locates the files on a LUN that is assigned to the SQL server for the `tempdb` file.

Create an additional `tempdb` file for each CPU core in each server. The Hitachi Data Systems test environment had 12 CPU cores on the test server.

For the implementation described in this installation guide, the following parameters shown in red replaced the placeholder variables:

```
USE tempdb
GO

ALTER DATABASE tempdb
ADD FILE (NAME='tempdev1', FILENAME= 'C:\SQLDBMNT\TempDB1\tempdb1.ndf ',
SIZE=420MB, MAXSIZE=120GB FILEGROWTH=42MB)


ALTER DATABASE tempdb
ADD FILE (NAME='tempdev2', FILENAME= 'C:\SQLDBMNT\TempDB2\tempdb2.ndf ',
SIZE=420MB, MAXSIZE=120GB FILEGROWTH=42MB)


ALTER DATABASE tempdb
ADD FILE (NAME='tempdev3', FILENAME= 'C:\SQLDBMNT\TempDB3\tempdb3.ndf ',
SIZE=420MB, MAXSIZE=120GB FILEGROWTH=42MB)
```

The procedure described in this section can be used also to relocate the Microsoft 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 Microsoft SQL Server for SharePoint 2010 deployments, see the Microsoft *TechNet* article "SQL Server and storage (SharePoint 2010)."

Install and Configure Microsoft SharePoint Server 2010

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

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

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

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