HITACHI
Inspire the Next

Microsoft SharePoint 2010 with Hitachi RBS Provider Adapter

Implementation Guide

Mark Crowley

January 2012

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 2010 with Hitachi RBS Provider Adapter

Implementation Guide

Use these instructions to implement a 200,000-seat Microsoft SharePoint 2010 farm, with Remote BLOB Storage (RBS) for the twenty content databases in the farm. The BLOB store volumes are SAN-attached volumes on Hitachi Virtual Storage Platform using economical SATA volumes.

This document uses Hitachi Data Systems lab naming conventions. Adapt all object naming (SAN zones, machine names, LUN names, and CIFS share names) to fit the deployment environment.

Tested Solution Components

These are the major components used in this tested solution.

Hardware Components

Table 1 describes the hardware resources used in the SharePoint farm.

Table 1. Hardware Resources

Hardware	*Description*	*Qty.*
Hitachi Virtual Storage Platform	4 × 8 GB Fibre Channel ports 69 GB cache memory 376 × 300 GB 10k RPM SAS drives 64 × 2 TB 7.2K RPM SATA drives	1
Hitachi Compute Blade 2000 chassis	2 × 1/10 Gb/sec Ethernet switch modules	1
Hitachi E55-A2 Server Blade	2 × Intel X5670 core processors 72 GB RAM 2 × 300 GB 10k RPM SAS hard drives Hitachi 4 Port 1 G Ethernet Mezzanine card 2 × Hitachi GVX-CC2N8G2X1 2-port PCIe HBA Fibre Channel cards	4
Brocade 5340 SAN Switch	80 8 Gb/sec Fibre Channel ports	2

The physical server configuration of the SharePoint farm is in Figure 1. It shows these components used in the Hitachi Data Systems lab for the following:

- Four server blades
- Virtual Storage Platform
- 1/10 Gb/sec switches
- Fibre Channel switches.

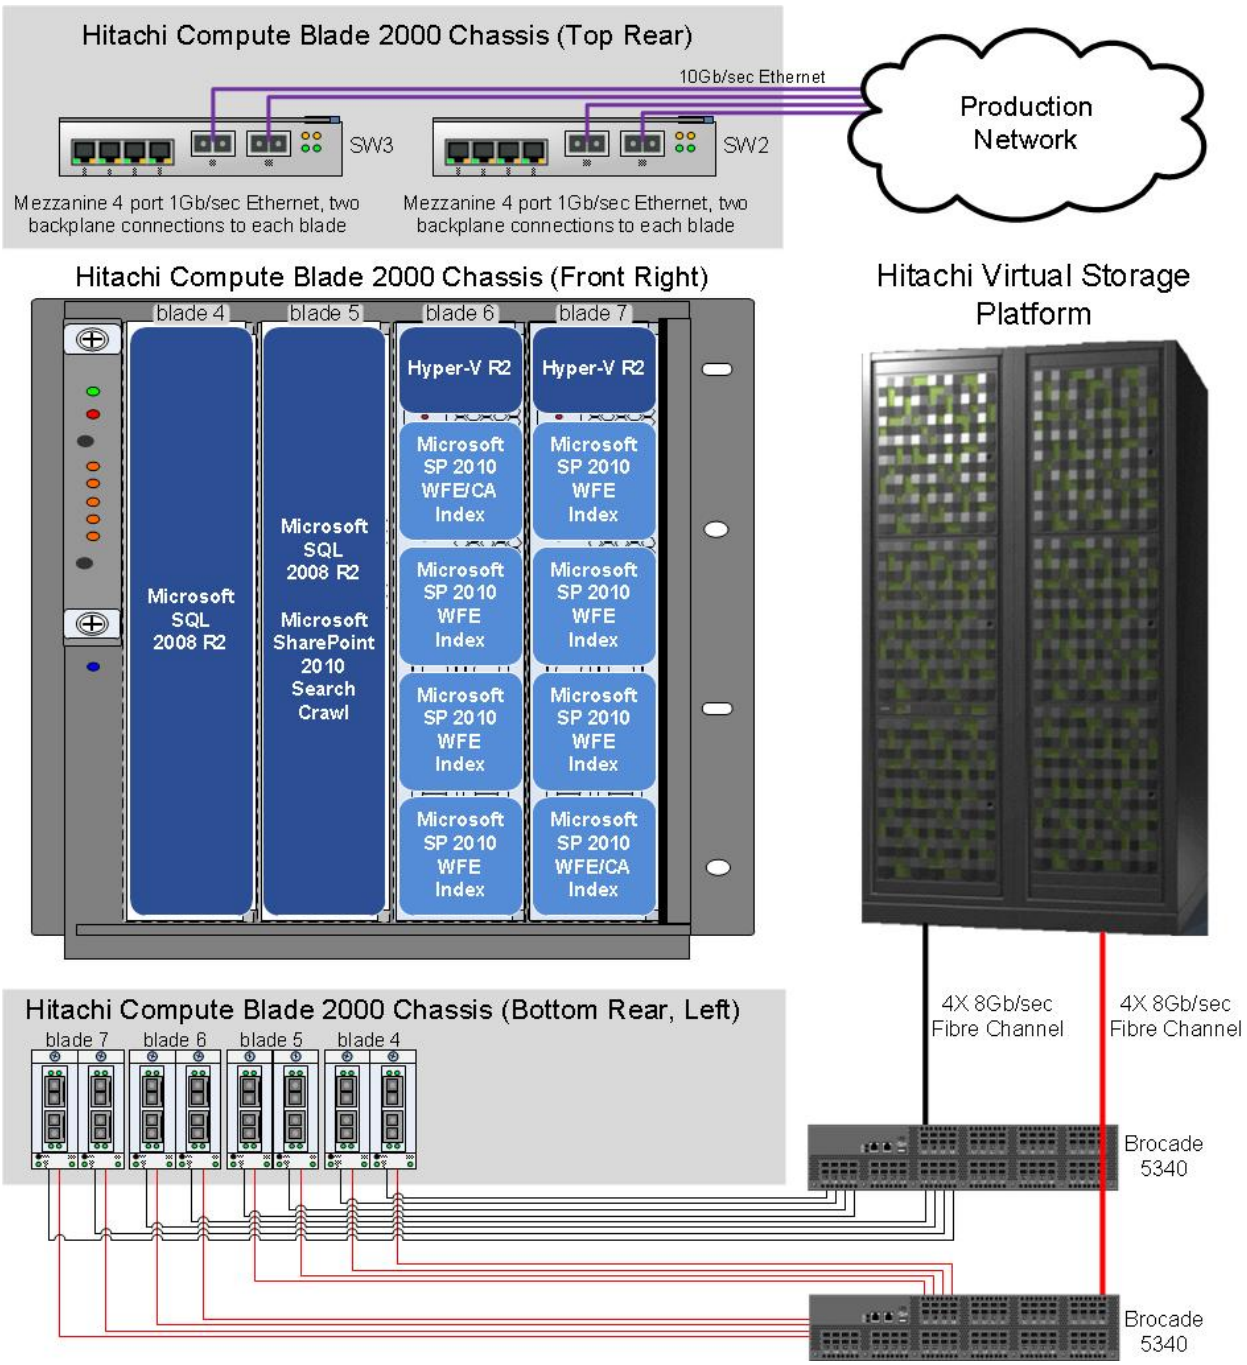

Figure 1

Servers of comparable speed and capacity can be used for this SharePoint farm. See the server hardware details in Table 1.

Software Components

Table 2 shows the operating system and applications installed on each server blade.

Table 2. Server and Role Details for the Server Blades

Hardware	*Virtual Machine*	*OS Installed*	*Applications and Roles Installed*
Blade 4	N/A	Microsoft Windows 2008 R2 Enterprise SP1	Microsoft SQL Server 2008 R2 Enterprise
Blade 5	N/A	Microsoft Windows 2008 R2 Enterprise SP1	Microsoft SharePoint 2010 Microsoft SQL Server 2008 R2 Enterprise
Blade 6	N/A	Microsoft Windows 2008 R2 Datacenter SP1	Microsoft Windows 2008 R2 Hyper-V
	SP-WS01	Microsoft Windows 2008 R2 Enterprise SP1	Microsoft SharePoint 2010
	SP-WS02	Microsoft Windows 2008 R2 Enterprise SP1	Microsoft SharePoint 2010
	SP-WS03	Microsoft Windows 2008 R2 Enterprise SP1	Microsoft SharePoint 2010
	SP-WS04	Microsoft Windows 2008 R2 Enterprise SP1	Microsoft SharePoint 2010
Blade 7	N/A	Microsoft Windows 2008 R2 Datacenter SP1	Microsoft Windows 2008 R2 Hyper-V
	SP-WS05	Microsoft Windows 2008 R2 Enterprise SP1	Microsoft SharePoint 2010
	SP-WS06	Microsoft Windows 2008 R2 Enterprise SP1	Microsoft SharePoint 2010
	SP-WS07	Microsoft Windows 2008 R2 Enterprise SP1	Microsoft SharePoint 2010
	SP-WS08	Microsoft Windows 2008 R2 Enterprise SP1	Microsoft SharePoint 2010

Note — Microsoft SharePoint 2010 requires the support of an Active Directory environment. In production environments, join the servers to the existing Active Directory domain before the installation of any components of Microsoft SQL Server and Microsoft SharePoint. The architecture used in the Hitachi Data Systems lab included an existing Active Directory installation.

Solution Implementation

These instructions assume that hardware has been previously installed and connected.

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 physical 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.
8. Create CIFS BLOB stores on the SharePoint client machines (web front ends)
9. Install RBS Components
10. Create CIFS BLOB Stores
11. Enable RBS

These are the general tasks necessary for a successful deployment. The following text gives details for each step. Your task list might vary from these steps, based on your environment and needs.

For more information about each of these tasks, see additional 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"
  - "Overview of RBS (SharePoint Server 2010)"

Configure Fabric Switch Zones

Configure zones on your fabric switches according to the manufacturer's guidelines, following these recommended 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 3 lists the Fibre Channel fabric zoning details used in the Hitachi Data Systems lab environment. Server names and Switch Zone names should be altered for the standards in the environment where this farm is to be deployed.

Table 3. 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_B4_HBA1_1_ASE45_36_3A	3A	5300-01
MSFT-SP-SQL-1	HBA1-2	BS2K_01_B4_HBA2_1_ASE45_36_4A	4A	5300-01
MSFT-SP-SQL-1	HBA2-1	BS2K_01_B4_HBA1_2_ASE45_36_7A	7A	5300-02
MSFT-SP-SQL-1	HBA2-2	BS2K_01_B4_HBA2_2_ASE45_36_8A	8A	5300-02
MSFT-SP-WS-SC	HBA1-1	BS2K_01_B5_HBA1_1_ASE45_36_3A	3A	5300-01
MSFT-SP-WS-SC	HBA1-2	BS2K_01_B5_HBA1_2_ASE45_36_4A	4A	5300-01
MSFT-SP-WS-SC	HBA2-1	BS2K_01_B5_HBA2_1_ASE45_36_7A	7A	5300-02
MSFT-SP-WS-SC	HBA2-2	BS2K_01_B5_HBA2_2_ASE45_36_8A	8A	5300-02
SANBOOT MSFT-SP-SQL-1	HBA1-1	BS2K_01_B4_HBA1_1_ASE45_36_4B	4B	5300-01
SANBOOT MSFT-SP-WS-SC	HBA1-1	BS2K_01_B5_HBA1_1_ASE45_36_4B	4B	5300-01

Configure the Hitachi Virtual Storage Platform

This is the storage architecture created for the Microsoft SharePoint environment using RBS. It takes into consideration Hitachi Data Systems and Microsoft recommended practices for the deployment of large SharePoint environments.

This document uses **Hitachi Storage Navigator** as the interface to manage the Hitachi Virtual Storage Platform. The specific configuration dictates pool sizing and arrangement of available parity groups on Virtual Storage Platform.

Dynamic Provisioning Pool Details

This SharePoint Farm uses four dynamic provisioning pools on Hitachi Virtual Storage Platform. These pools are composed of LDEVs which are created when initially defining the parity groups.

Table 4 provides details on the pool configurations.

Table 4. Dynamic Provisioning Pools

Dynamic Provisioning Pool Name (ID)	*RAID Group Config.*	*Drive Type*	*Number of RAID Groups*	*Pool Capacity (TB)*
DB ▪ SQL database files (.mdf and .ndf) ▪ tempDB database files (.mdf) ▪ Dynamic provisioning pool with Hitachi Dynamic Tiering enabled	RAID-5 (7D+1P)	300 GB 10k RPM SAS ▪ Tier 1 storage for frequently accessed data	1	7.00
	RAID-5 (7D+1P)	2 TB 7.2k RPM SATA ▪ Tier 2 storage for remaining data, using more cost-effective drives	2	
LOG ▪ SQL database transaction log files (.ldf) ▪ tempDB transaction log files (.ldf) ▪ Dynamic provision pool	RAID-5 (7D+1P)	300 GB 10k RPM SAS	1	1.87
VMPOOL ▪ Guest virtual machines boot and index volumes ▪ Dynamic provisioning pool	RAID-5 (7D+1P)	300 GB 10k RPM SAS	1	1.87
BLOBSTORE ▪ BLOB stores that are presented to the SharePoint virtual machines. ▪ Dynamic provisioning pool	RAID-5 (7D+1P)	2 TB 7.2k RPM SATA	2	5.13

Launch Hitachi Storage Navigator

To launch Hitachi Storage Navigator, do the following:

1. Open Microsoft Internet Explorer.
2. Type the following in the address bar:

   `http://<SVP IP ADDRESS>/`

3. Log on to the SVP console.

The Hitachi Storage Navigator console opens.

The links to the wizards referred to in this guide are located on the **General Tasks** pane in the main window.

Create Dynamic Provisioning Pools

Follow these steps to create the pools:

1. Open the Create Pools Wizard.

   (1) From the **Explorer** panel, click **Pools**. The right panel shows previously configured pools on the storage system.

   (2) At the bottom of that window, click **Create Pools**.

2. Add one or more LDEVs to a dynamic provisioning pool.

   (1) From the **Pool Type** list, click **Dynamic Provisioning**.

   (2) (optional) Click **Enable** for the **Multi-Tier Pool** option (database pool).

   (3) Click **Select Pool VOLs**. The **Select Pool VOLs** dialog box opens.

   (4) In the new dialog box, click the heading **Drive Type/RPM** to sort the LDEVS.

   (5) Select one or more LDEV to add them to the pool.

   If the pool is a **Multi-Tier Pool**, the Virtual Storage Platform marks each selected LDEV with a tier number. The fastest LDEVs are marked as Tier 1.

   i. **Top Tier** — Select an LDEV from the fastest **LDEVs** available by **Drive Type/RPM**, and then click **ADD**. The selected **LDEV** moves to the right panel.

   This positions this LDEV as the top volume in the pool.

   ii. **Other Tier** — Select the remaining LDEV or LDEVs to be added to the pool and then click **ADD**.

   (6) When done adding LDEVs, click **OK**.

   The **Select Pool VOLs** window closes. Return to the Create Pools Wizard.

   (7) Check the quantity of the pool volumes and pool size numbers under **Select Pool VOLs**.

3. Name and set other details for the dynamic provisioning pool.

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

   The Hitachi Data Systems lab environment uses the prefixes in Table 5.

   Table 5. Dynamic Provisioning Pool Name Prefixes

Type	*Prefix*
Database pool	SP-DB
Log pool	SP-LOG
Virtual machine storage	SP-VMPOOL
BLOB stores	SP-BLOBSTORE

   (2) (Optional) Type the initial number for the first pool name in **Initial Number**.

   (3) Expand the **Options** area.

(4) Type the 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 shows each pool that is being created.

(7) Click **Finish**.

The **Create Pools** dialog box opens.

(8) (Optional) From the **Create Pools** dialog box, click **Next** 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 "Create Host Groups."

(9) Click **Apply**.

Repeat these steps to create the other dynamic provisioning pools.

SharePoint LDEV Details

The following tables detail the LDEVs created within the pools for the Hitachi Data Systems LAB SharePoint farm.

The pool names, the pool indices (the number in parenthesis in the pool name column), the LDEV numbers, and the storage ports are specific to the Hitachi Data Systems lab SharePoint farm, and should be altered to fit any other SharePoint installation.

The volume sizes are the recommended sizes for a 200,000 seat SharePoint farm.

All LDEVs created from the various pools are thin provisioned on Hitachi Virtual Storage Platform. In practice, the actual space usage for each LDEV is reduced significantly from the maximum capacity of the LDEV, especially for the database volumes. These numbers represent the maximum space available for each component.

Table 6 lists the volumes provisioned for the configuration database.

Table 6. Configuration Database Volumes

Dynamic Provisioning pool Name (ID)	*Mapped to Server*	*LDEV*	*Size (GB)*	*Purpose*	*Storage Port*
SP-DB (30)	MSFT-SP-SQL-1	30:00	10	Database	3A/4A/7A/8A
SP-LOG (31)	MSFT-SP-SQL-1	31:00	1	Log	3A/4A/7A/8A

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

Table 7. Central Administration Database Volumes

Dynamic Provisioning Pool Name (ID)	*Mapped to Server*	*LDEV*	*Size (GB)*	*Purpose*	*Storage Port*
SP-DB (30)	MSFT-SP-SQL-1	30:01	10	Database	3A/4A/7A/8A
SP-LOG (31)	MSFT-SP-SQL-1	31:01	1	Log	3A/4A/7A/8A

Table 8 lists the volumes provisioned for the content databases.

Table 8. Content Databases Volumes

Dynamic Provisioning Pool Name (ID)	*Mapped to Server*	*LDEVs*	*Size (GB)*	*Purpose*	*Storage Port*
SP-DB (30)	MSFT-SP-SQL-1	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	SharePoint Content Databases 00-19	3A/4A/7A/8A
SP-LOG (31)	MSFT-SP-SQL-1	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	SharePoint Content Database Logs 00-19	3A/4A/7A/8A

Table 9 lists the volumes provisioned for the search administration database.

Table 9. Search Administration Volumes

Dynamic Provisioning Pool Name (ID)	Mapped to Server	LDEV	Size (GB)	Purpose	Storage Port
SP-DB (30)	MSFT-SP-SQL-1	30:16	10	Database	3A/4A/7A/8A
SP-LOG (31)	MSFT-SP-SQL-1	31:16	2	Log	3A/4A/7A/8A

Table 10 lists the volumes provisioned for the crawl database.

Table 10. Crawl Database Volumes

Dynamic Provisioning Pool Name (ID)	Mapped to Server	LDEV	Size (GB)	Purpose	Storage Port
SP-DB (30)	MSFT-SP-WS-SC	30:17	185	Database	3A/4A/7A/8A
SP-LOG (31)	MSFT-SP-WS-SC	31:17	36	Log	3A/4A/7A/8A

Table 11 lists the volumes provisioned for the search property database.

Table 11. Search Property Database Volumes

Dynamic Provisioning Pool Name (ID)	Mapped to Server	LDEV	Size (GB)	Purpose	Storage Port
SP-DB (30)	MSFT-SP-WS-SC	30:18	60	Database	3A/4A/7A/8A
SP-LOG (31)	MSFT-SP-WS-SC	31:18	15	Log	3A/4A/7A/8A

Table 11 lists the volumes provisioned for the `tempdb` files.

Table 12. `tempDB` Volumes

Dynamic Provisioning Pool Name (ID)	Mapped to Server	LDEV	Size (GB)	Purpose	Storage Port
SP-DB (30)	MSFT-SP-SQL-1	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	TempDB Databases 0-11	3A/4A/7A/8A
SP-LOG (31)		30:31	320	Log	3A/4A/7A/8A

Table 13 lists the volumes provisioned for the Hyper-V virtual machines and Index volumes.

Table 13. Virtual Machine Volumes

Dynamic Provisioning Pool Name (ID)	LDEV	Size (GB)	Mapped to VMhost -> Virtual machine	Storage Port
SP-VMPOOL (35)	30:20	150	BLADE6-HV -> SP-WS01 boot volume	3B/7B
SP-VMPOOL (35)	30:24	150	BLADE6-HV -> SP-WS01 index volume	3B/7B
SP-VMPOOL (35)	30:21	150	BLADE6-HV -> SP-WS02 boot volume	3B/7B
SP-VMPOOL (35)	30:25	150	BLADE6-HV -> SP-WS02 index volume	3B/7B
SP-VMPOOL (35)	30:22	150	BLADE6-HV -> SP-WS03 boot volume	3B/7B
SP-VMPOOL (35)	30:26	150	BLADE6-HV -> SP-WS03 index volume	3B/7B
SP-VMPOOL (35)	30:23	150	BLADE6-HV -> SP-WS04 boot volume	3B/7B
SP-VMPOOL (35)	30:27	150	BLADE6-HV -> SP-WS04 index volume	3B/7B
SP-VMPOOL (35)	30:28	150	BLADE7-HV -> SP-WS05 boot volume	4B/8B
SP-VMPOOL (35)	30:2C	150	BLADE7-HV -> SP-WS05 index volume	4B/8B
SP-VMPOOL (35)	30:29	150	BLADE7-HV -> SP-WS06 boot volume	4B/8B
SP-VMPOOL (35)	30:2D	150	BLADE7-HV -> SP-WS06 index volume	4B/8B
SP-VMPOOL (35)	30:2A	150	BLADE7-HV -> SP-WS07 boot volume	4B/8B
SP-VMPOOL (35)	30:2E	150	BLADE7-HV -> SP-WS07 index volume	4B/8B
SP-VMPOOL (35)	30:2B	150	BLADE7-HV -> SP-WS08 boot volume	4B/8B
SP-VMPOOL (35)	30:2F	150	BLADE7-HV -> SP-WS08 index volume	4B/8B

Table 14 lists the volumes used for the BLOB store shares.

Table 14. BLOB Store Volumes

Dynamic Provisioning Pool Name (ID)	LDEV	Size (GB)	Mapped to VMhost -> Virtual machine	Storage Port
SP-BLOBSTORE (32)	30:00	200	BLADE6-HV -> SP-WS01 RBSCDB01	3B/7B
SP-BLOBSTORE (32)	30:01	200	BLADE6-HV -> SP-WS01 RBSCDB02	3B/7B
SP-BLOBSTORE (32)	30:02	200	BLADE6-HV -> SP-WS02 RBSCDB03	3B/7B
SP-BLOBSTORE (32)	30:03	200	BLADE6-HV -> SP-WS02 RBSCDB04	3B/7B
SP-BLOBSTORE (32)	30:04	200	BLADE6-HV -> SP-WS02 RBSCDB05	3B/7B
SP-BLOBSTORE (32)	30:05	200	BLADE6-HV -> SP-WS03 RBSCDB06	3B/7B
SP-BLOBSTORE (32)	30:06	200	BLADE6-HV -> SP-WS03 RBSCDB07	3B/7B
SP-BLOBSTORE (32)	30:07	200	BLADE6-HV -> SP-WS04 RBSCDB08	3B/7B
SP-BLOBSTORE (32)	30:08	200	BLADE6-HV -> SP-WS04 RBSCDB09	3B/7B
SP-BLOBSTORE (32)	30:09	200	BLADE6-HV -> SP-WS04 RBSCDB10	3B/7B
SP-BLOBSTORE (32)	30:0A	200	BLADE7-HV -> SP-WS05 RBSCDB11	4B/8B
SP-BLOBSTORE (32)	30:0B	200	BLADE7-HV -> SP-WS05 RBSCDB12	4B/8B
SP-BLOBSTORE (32)	30:0C	200	BLADE7-HV -> SP-WS06 RBSCDB13	4B/8B
SP-BLOBSTORE (32)	30:0D	200	BLADE7-HV -> SP-WS06 RBSCDB14	4B/8B
SP-BLOBSTORE (32)	30:0E	200	BLADE7-HV -> SP-WS06 RBSCDB15	4B/8B

Dynamic Provisioning Pool Name (ID)	LDEV	Size (GB)	Mapped to VMhost -> Virtual machine	Storage Port
SP-BLOBSTORE (32)	30:0F	200	BLADE7-HV -> SP-WS07 RBSCDB16	4B/8B
SP-BLOBSTORE (32)	30:10	200	BLADE7-HV -> SP-WS07 RBSCDB17	4B/8B
SP-BLOBSTORE (32)	30:11	200	BLADE7-HV -> SP-WS08 RBSCDB18	4B/8B
SP-BLOBSTORE (32)	30:12	200	BLADE7-HV -> SP-WS08 RBSCDB19	4B/8B
SP-BLOBSTORE (32)	30:13	200	BLADE7-HV -> SP-WS08 RBSCDB20	4B/8B

Create LDEVs within the Pools

To create a LDEV, follow these steps:

1. Start creating the LDEV or LDEVs.
   (1) On the **Explorer** panel, select **Pools.**
   (2) Click the pool name in which to create one or more LDEVs.
   (3) Click **Create LDEVs**.
2. Enter the size of each LDEV and number of LDEVs to create.
   (1) Type a size for each LDEV to be created in **LDEV Capacity**.
   (2) Type the number of LDEVs to be created in **Number of LDEVs**.
3. Enter the prefix and starting number for the LDEVs.
   (1) Type a prefix for the LDEV or LDEVs that you are creating in **Prefix**.
   Use something descriptive, like **SP-CDB-**.
   (2) Type the starting number in **Initial Number** using two digits when creating more than one LDEV (for example, **01**).
   This number increases by 1 for each LDEV created.
4. Set the pool number and starting number of the LDEV.
   (1) Expand the **Options** section under the **LDEV Name** section.
   (2) 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.
   (3) 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.
5. Click **Add**. The created LDEV or LDEVs populate the **Selected LDEVs** column.
6. Click **Finish,** and then click **Apply**.

Create Host Groups

Enable port security on the Hitachi Virtual Server Platform ports used for this solution before creating host groups.

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

1. Start the Create Host Groups Wizard.

   (1) From the **Explorer** panel, click **Ports/Host Groups**.

   (2) At the bottom of the **Storage Systems** panel, click **Create Host Groups**.

   The **Create Host Groups** wizard opens.

2. Type a name for the group in **Host Group Name**.
3. From the **Host Mode** list, click **2C [Windows Extended]**.
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**.

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. Write down the **LDEV ID** and **LUN ID** correlations from this window. This information will be used later.

9. Click **Finish**.

   The **Add LUN Paths** window opens.

10. Click **Apply**.

Configure the Physical Servers

Install the appropriate edition of Microsoft Windows Server 2008 R2 on the hosts. Refer to Figure 1 and Table 2 for the operating system versions used on the 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, such as the HBA management software.

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

Set up mount points on the physical hosts

See the tables in "SharePoint LDEV Details" the mappings used when validating this reference architecture for a SharePoint Farm. Use appropriate names for your network architecture.

To map the presented volumes to mount points on the servers, do the following:

1. On the SQL server, open **Windows Explorer**, and then click **Local Disk (C:)**

2. Create the follow folders:

   (1) In the root directory of drive C, create these folders:

   - C:\SQLData
   - C:\SQLLog
   - C:\SQLSystem

   (2) In **C:\SQLData**, create folders for the following SharePoint mount points:

   - DB LUNs
   - TempDB LUNs

   Example folders include C:\SQLData\SPCDB01 and C:\SQLData\TEMPDB01.

(3) In **C:\SQLLog**, create folders for the following SharePoint mount points:

- Log LUNs
- TempDB log LUN

Example folders include C:\SQLLog\SPCLOG01 and C:\SQLLog\TEMPLOG.

(4) In **C:\SQLSystem**, create folders for the following SharePoint mount points:

- Central Administration and log LUNs
- Admin Content and log LUNs
- Search Administration DB and LOG LUNs

Example folders include C:\SQLSystem\SPCADB and C:\SQLSystem\SPCALOG.

3. Initialize Disk.

(1) Open **Server Manager**.

(2) In the left pane, expand the **Storage** section, and left click **Disk Management**.

The central pane in the window populates with information about all of the volumes attached to the system. The new LUNs are marked **Offline**.

- If the LUNs are not present, right click **Disk Management** and then click **Rescan Disks.** This refreshes the list of volumes.

The bottom half of the central pane displays the actual disk numbers.

(3) Right-click the heading (the leftmost section of each item) for the first offline disk, and then click **Properties**.

A new dialog box opens.

The LUN number in the **Location** line is the actual LUN ID of the volume from the list of LUN IDs that was recorded when the LUNs were mapped. (See “Map LDEVs to Host Groups.")

(4) Close the **Hitachi Open-V Disk Device Properties** window.

(5) Right-click the heading of the same disk, and click **Online**.

The status of the disk changes to **Online**. Right-click again, and click **Initialize Disk**.

4. Format Volume.

(1) Right click the **Information** section for the selected disk, and then click **New Simple Volume…**

The New Simple Volume Wizard opens.

(2) Click **Next** and then **Next**.

(3) On the **Assign Drive Letter or Path** page of the wizard, Click **Mount in the following empty NTFS folder**, and then click **Browse.**

A selection window opens.

(4) Click the folder for this LUN number.

The path to that location populates the field to the left of the **Browse** button.

(5) Click **Next**.

(6) Click **Format this volume with the following settings** and then set the parameters from Table 15:

Table 15. Disk Volume Settings

Setting	*Parameter*
File system	NTFS
Allocation unit size	64 KB The 64 KB allocation unit size is recommended for Microsoft SQL Server volumes.
Volume Label	Type the volume label.

(7) Verify that the **Perform a quick format** checkbox is selected, and then click **Next**.

(8) Click **Finish**.

The window closes, and the LUN formats and mounts.

5. Create a folder in the root directory of the mount point. For example, create the directory SQLDATA. This folder is for storing SQL files.

Repeat these steps for every LUN on the system.

Repeat these steps on the SharePoint Search Crawl server for consistency.

Note – The LUNs mapped to the Hyper-V hosts are used as pass-through disks. Hyper-V assigns these to the virtual machines in "Enable the Hyper-V Role on the Virtual Hosts."

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 **Add Roles Wizard**.

   The Add Roles wizard opens.
3. Click **Next**.

   A list of available roles displays.
4. Select the **Hyper-V** check box, and then click **Next**.
5. Answer the questions for the system's configuration options, and **do not select** any network adapters. Selecting adapters adds unnecessary complexity to the network configuration done later. 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.

6. Run **Windows Update** on the machine.

7. From the **Administrative tools**, open the **Hyper-V Manager**. Click on the **Virtual Network Manager** option in the **Actions** panel. The **Virtual Network Manager** dialog box opens.

Hyper-V uses virtual network switches. These may have a dedicated connection to a physical network interface. From that, virtual NICs (VNIC) may be assigned to the host system and to virtual machines. The specific configuration is dependant on the network configuration for the individual host and available network adapters.

In the lab environment, each of the two Hyper-V hosts contained six 1G Ethernet adapters, as follows:

- One adapter is dedicated to host access.
- Four adapters are available to the virtual machines.

Each virtual machine has two VNICs assigned to them, as follows:

- One for SQL server access
- One for Web services

All VNICs IP addresses are in a single /24 subnet.

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

You have various options to deploy the virtual machines within a Microsoft Hyper-V environment. Use the best method for your environment when deploying the virtual machines on the two Hyper-V enabled hosts.

In the testing lab farm, there are four virtual machines on each Hyper-V host, each connected to two Hyper-V virtual networks, as follows:

- One for SQL traffic
- One for web traffic (using Microsoft Network load balancing).

All disk volumes are pass-through disks.

For more information about setting up Microsoft Hyper-V to deploy virtual machines, see Hyper-V on Microsoft *TechNet.* Take special note of these sections:

- Getting Started
- Planning
- Installation
- Configuration

Run the **Hyper-V Best Practices Analyzer** to verify appropriate configuration for your environment. 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.

As part of the SQL 2008 R2 install, ensure that the following collation settings are selected to suit SharePoint database use:

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

You may use the default location for the database and log files at this time. After completing the installation, move the `tempdb` databases to 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

In the Hitachi Data Systems lab environment, the following are mapped:

- 83 LUNs mapped from the database and log pools to the primary SQL server
- 36 LUNs mapped from the VMPOOL and BLOBSTORE pools to the two Hyper-V hosts
- 4 LUNs mapped to the Search Crawl server.

Refer to the tables in "Create LDEVs within the Pools" and "Set up mount points on the physical hosts" before installing Microsoft SQL Server.

"Databases and Transaction Log Files" and "`tempdb` Files" 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, establish an allocation strategy using these configuration parameters:

- SIZE
- FILEGROWTH
- MAXSIZE

For the database, the values of SIZE, FILEGROWTH, and MAXSIZE are determined by expected growth.

- The minimum file size for all of the databases was changed to 420 MB. The minimum file size for the logs was changed to 42 MB. Use a multiple of 42 MB when using Hitachi Dynamic Provisioning and Hitachi Dynamic Tiering because the page size is 42 MB.
- The file auto extends in size by the amount specified in FILEGROWTH when the currently allocated space for the file runs out. Use a value for FILEGROWTH that is a multiple of 42 MB.
- The file growth stops if it reaches the amount specified in MAXSIZE. This keeps your application from exceeding the 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 in many small increments. This can slow database startup as well as the 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 42 MB for implementations that use Hitachi Dynamic Provisioning or Hitachi Dynamic Tiering.

Use the TSQL code samples in Figure 2 and Figure 3 to create the required databases on your SQL Server installation.

Execute the SQL statements in Figure 2 using the query window opened by the **New Query** option in Microsoft 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>)
```

Figure 2

For the implementation described in this installation guide, the parameters in Figure 3 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)
```

Figure 3

Use the instant file initialization feature of Microsoft SQL Server for 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 (create zero pages). Instead, disk content is overwritten as new data is written to the files. This Microsoft SQL Server feature works in conjunction with Hitachi Dynamic Provisioning. For more information, see Database File Initialization on SQL Server 2008 Books Online.

tempdb Files

The default location for tempdb files in Microsoft 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 Hitachi Virtual Storage Platform. To do this, execute the SQL statements in Figure 4 using 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>)
```

Figure 4

For the implementation described in this installation guide, the parameters in Figure 5 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)
```

Figure 5

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 parameters shown in red in Figure 6 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)
```

Figure 6

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

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."

Create CIFS BLOB Stores on the SharePoint Client Machines (Web Front Ends)

Map the LDEVs that are intended to contain the Blob stores to the Hyper-V hosts that contain the virtual machines serving the SharePoint web front-end role. In the environment referenced in this document, there are 8 virtual machines running this role on two Hyper-V hosts. There are 20 LDEVs created for the 20 content databases (enabled databases for RBS), with 10 mapped to each Hyper-V host.

The volumes are distributed in the lab environment as pass-through disks to the four virtual machines on each Hyper-V host. Table 16 details how the BLOB store pass-through disks were mapped.

Table 16. Blob Store LUN Maps and Share Details

Hyper-V Host	*Virtual Machine Name*	*BLOB Store Name*	*CIFS Share*
BLADE 6-HV	SP-WS01	RBSCDB01 RBSCDB02	\\SP-WS01\RBSCDB01 \\SP-WS01\RBSCDB02
BLADE 6-HV	SP-WS02	RBSCDB03 RBSCDB04 RBSCDB05	\\SP-WS02\RBSCDB03 \\SP-WS02\RBSCDB04 \\SP-WS02\RBSCDB05
BLADE 6-HV	SP-WS03	RBSCDB06 RBSCDB07	\\SP-WS03\RBSCDB06 \\SP-WS03\RBSCDB07
BLADE 6-HV	SP-WS04	RBSCDB08 RBSCDB09 RBSCDB10	\\SP-WS04\RBSCDB08 \\SP-WS04\RBSCDB09 \\SP-WS04\RBSCDB10
BLADE 7-HV	SP-WS05	RBSCDB11 RBSCDB12	\\SP-WS05\RBSCDB11 \\SP-WS05\RBSCDB12
BLADE 7-HV	SP-WS06	RBSCDB13 RBSCDB14 RBSCDB15	\\SP-WS06\RBSCDB13 \\SP-WS06\RBSCDB14 \\SP-WS06\RBSCDB15
BLADE 7-HV	SP-WS07	RBSCDB16 RBSCDB17	\\SP-WS07\RBSCDB16 \\SP-WS07\RBSCDB17
BLADE 7-HV	SP-WS08	RBSCDB18 RBSCDB19 RBSCDB20	\\SP-WS08\RBSCDB18 \\SP-WS08\RBSCDB19 \\SP-WS08\RBSCDB20

Volumes on each virtual machine (SP-WS0X) are mapped as UNC/CIFS shares using “Advanced Sharing.” Full control was granted for the SharePoint farm account. For example, the RBSCDB01 share (the BLOB store for SPCDB01) is mapped as \\SP-WS01\RBSCDB01 for the farm.

The naming in the laboratory environment SharePoint 2010 farm are provided as examples to show the structure and distribution of LUNs to the various machines. Use naming appropriate for your environment.

Install RBS Components

The RBS infrastructure is comprised of two different installation packages, the *Microsoft RBS Feature Pack for SQL server 2008 (X64 version)*, and the RBS Provider Adapter for Hitachi Storage.

The Microsoft RBS Feature Pack install creates the appropriate tables and stored procedures within a specified database to support RBS. Copy the packages to a location, such as drive C, on every server in the farm.

Prepare SQL databases for RBS

Enable master key encryption in the content databases. Use the TSQL statement in Figure 7 using Microsoft SQL Server 2008 R2 Management Studio.

```
USE [<YourDatabaseName>]
CREATE MASTER KEY ENCRYPTION BY PASSWORD = '<YourPassword>';
GO
```

Figure 7

Execute this TSQL command against every database that supports RBS. (For SharePoint 2010, execute this against content databases only. It is not necessary to execute this against other SharePoint 2010 databases.)

Install the Microsoft RBS Feature Pack on the SQL server

To install the Microsoft RBS Feature Pack on the SQL server, do the following:

1. Open a command prompt on the SQL server
2. Log on to the SQL server as the database owner, typically the SharePoint farm account.
3. Navigate to the folder where RBS.msi resides.
4. Execute the command in Figure 8.to install RBS features to the database.

```
msiexec /lvx* \LOGS\rbs_enable_SPCDB01.log /q /i rbs.msi ADDLOCAL=EnableRBS
DBINSTANCE="<YourDatabaseServerName>" DBNAME="<YourDatabaseName>"
```

Figure 8

Execute this *msiexec* command on the SQL server once for every RBS-enabled content database. (Do not run this on the SharePoint administration databases.)

The log file name specified in the command (in this example, “\LOGS\rbs_enable_SPCDB01.log”) need to be changed for each execution to avoid overwriting the previous log file on the next execution.

Verify the install by looking in each log file. Check for the presence of several new tables written into the specified database with the “mssqlrbs_resources” name prefix.

Install the Microsoft RBS Feature Pack on the Client Machines

To install the Microsoft RBS Feature Pack on the web and application servers, do the following:

1. Open a command prompt.
2. Navigate to the path where RBS.MSI resides.
3. Execute the command in Figure 9 to install RBS binaries to the hosts.

```
msiexec /lvx* rbs_install.log /q /i rbs.msi
ADDLOCAL=ServerScript,Client,Maintainer
```

Figure 9

Repeat this process for every web or application server in your farm.

Install RBS Package on the Client Machines

Install Microsoft Windows® PowerShell Snap-in for Hitachi Storage on every web and application server. Included with RBS Provider for Hitachi Storage are two .msi files that require installation.

1. Execute the .msi file for Microsoft Windows PowerShell Snap-in for Hitachi Storage.

   ```
   Hitachi-pssnapin-setup(04.0.5)-(x64).msi
   ```

2. Execute the RBS Admin console installer.

   ```
   hitachi-rbsprovider-setup(01.2.1)-(x64).exe
   ```

   When this installer finishes, the icon to launch this program displays on the desktop.

Use RBS Provider for Hitachi Storage to Create and Register BLOB Stores

To create and register BLOB stores, do the following:

1. Open RBS Provider for Hitachi Storage. The console window opens, as shown in Figure 10

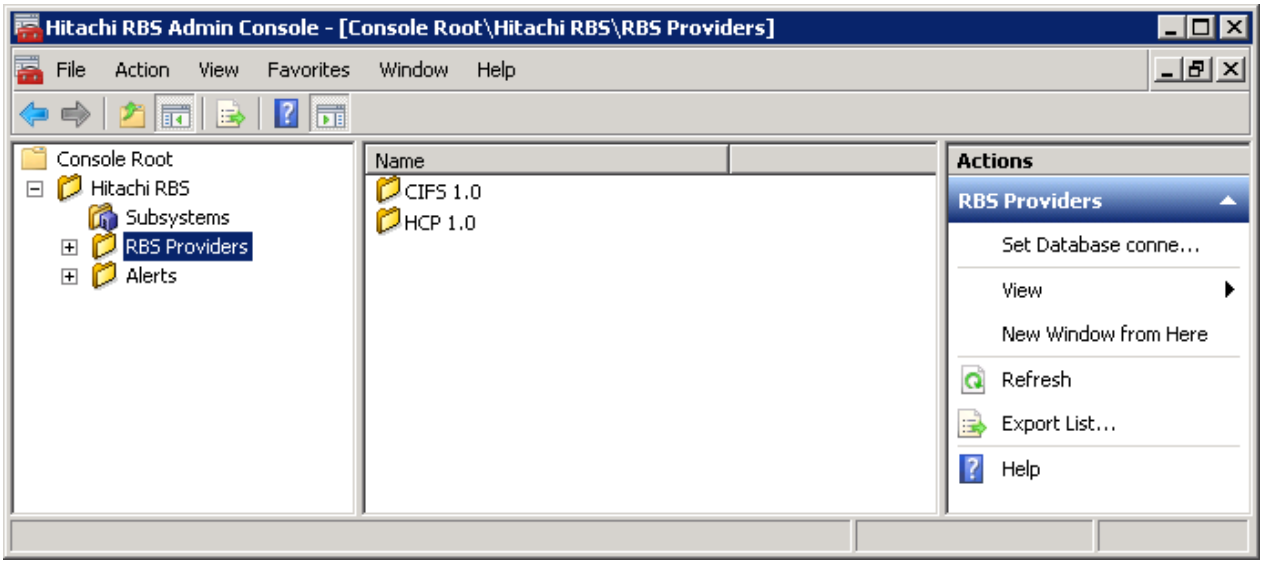

Figure 10

2. Click on the **RBS Providers** tree to expand the tree. You will use this later.

3. Right-click **RBS Providers** and click **Set Database connection data**. The dialog box in Figure 11 opens.

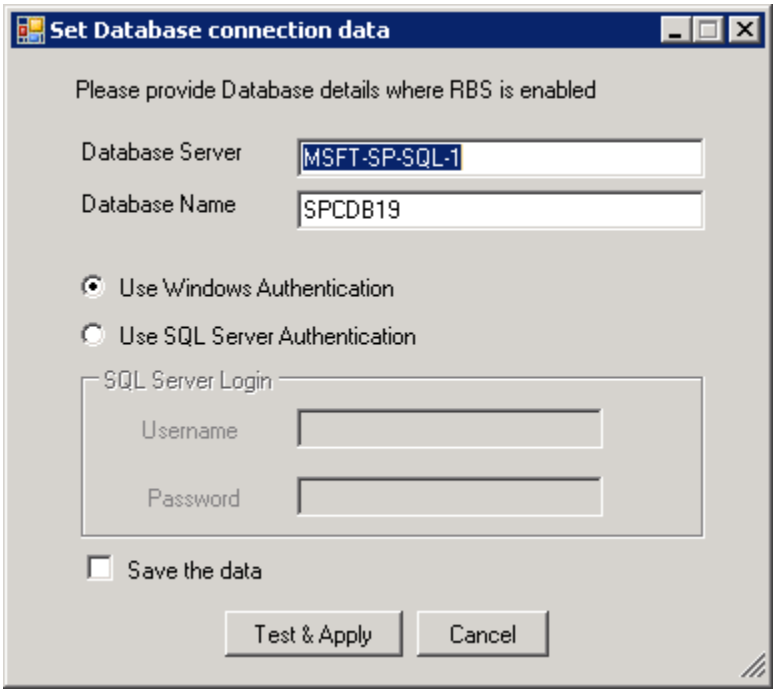

Figure 11

4. Type the database server name and the database name of the first database for which this system stores BLOBs.

   This process connects this Admin console to the named database. It inserts the BLOB store name into the RBS table in the database.

5. Click **Test & Apply**.

6. In tree under **RBS Providers,** select **CIFS 1.0.** See Figure 12.

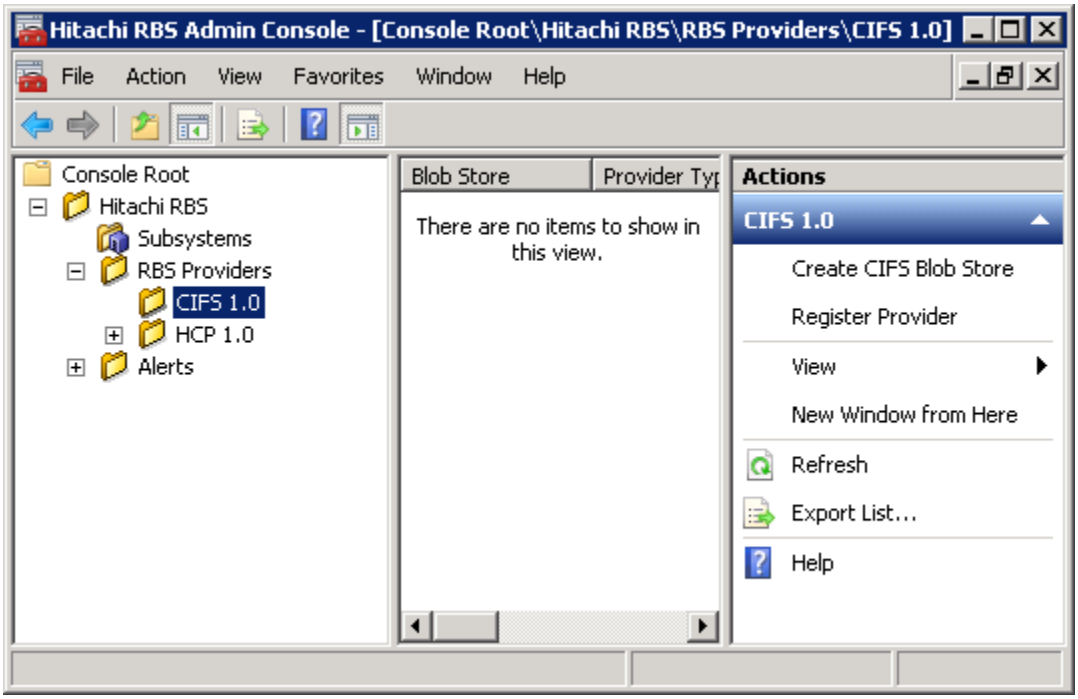

Figure 12

Create CIFS BLOB Stores

To create CIFS BLOB stores, do the following:

1. In the **Actions** pane, click **Create CIFS Blob Store**. A new dialog box opens.

2. Type the name in **BLOB Store Name** and 99999 in **BLOB Pool Capacity (#BLOBs)**. Select the **Set as Default** checkbox, and click **Next**.

3. Select the **Use UNC Path** checkbox, and type the path to the share that you created earlier to the BLOB store for this database. Click **Next**.

4. Clear the **Set Store Credentials** checkbox, and click **Next**.

5. Click **Create**. A completion dialog box confirms the completion of the BLOB store record in the database.

Repeat these steps for each content database on every web server.

When finished, there is an icon for **Launch Hitachi RBS Admin Console** on the desktop. It has details about the mapped RBS store to the last mapped content database in the middle panel.

Enable RBS

Enable RBS in each database after creating the BLOB store once in each content database. Use **SharePoint 2010 Management Shell**, which was installed with Microsoft SharePoint 2010.

When running SharePoint 2010 Management Shell, use the **Run as administrator** option.

Type the commands in Figure 13 at the PowerShell prompt:

```
$cdb = Get-SPContentDatabase <SharePoint database name>
$rbss = $cdb.RemoteBlobStorageSettings
$rbss.SetActiveProviderName($rbss.GetProviderNames()[0])
$rbss.Enable()
```

Figure 13

Type **$rbss <enter>** at the prompt to confirm that the setting is complete. Figure 14 shows the command flow and appropriate feedback with the SPCDB01 database connected to the RBSCDB01 BLOB store.

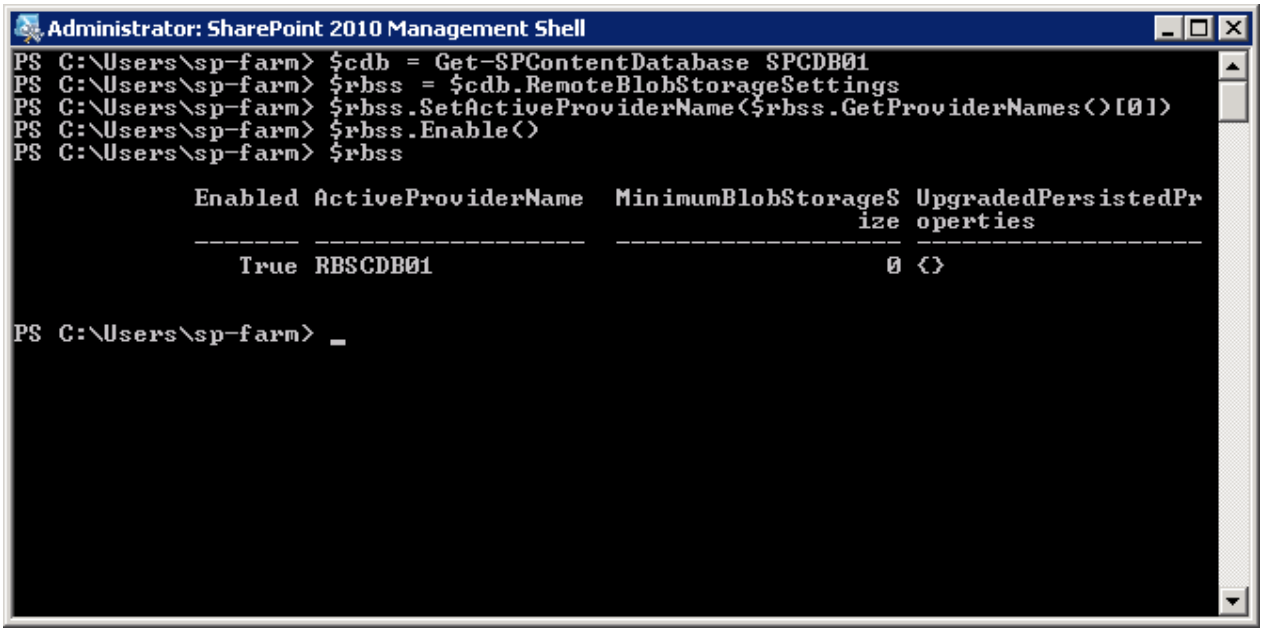

Figure 14

Repeat the PowerShell commands for all of your content databases.

Set up the RBS Maintainer Tasks

The RBS Maintainer executable requires connection strings manually entered into a configuration file to find the RBS databases. Run the maintainer through the Windows Task Scheduler.

The configuration file **Microsoft.Data.SqlRemoteBlobs.Maintainer.exe.config** installs with the Microsoft RBS Feature Pack installation in this location: **C:\Program Files\Microsoft SQL Remote Blob Storage 10.50\Maintainer**.

By default, the connection strings are encrypted in the default configuration file. Remove everything between <connectionstrings> and </connectionstrings>, as shown in Figure 15.

```
<configuration>
    <connectionStrings>
    </connectionStrings>
    <RemoteBlobStorage>
        <Logging>
            <add key="ConsoleLog" value="0" />
        </Logging>
    </RemoteBlobStorage>
</configuration>
```

Figure 15

Create a connection string for each content database that you have RBS enabled on, as shown in Figure 16.

```
<configuration>
    <connectionStrings>
   <add name="<first connection string name>" connectionString="Data
Source=<SQL server name>;Initial Catalog=<first content database
name>;Integrated Security=True;Application Name=&quot;Remote Blob Storage
Maintainer&quot;" providerName="System.Data.SqlClient" />
   <add name="<second connection string name>" connectionString="Data
Source=<SQL server name>;Initial Catalog=<second content database
name>;Integrated Security=True;Application Name=&quot;Remote Blob Storage
Maintainer&quot;" providerName="System.Data.SqlClient" />
   <add name="<third connection string name…>" connectionString="Data
Source=<SQL server name>;Initial Catalog=<second content database
name>;Integrated Security=True;Application Name=&quot;Remote Blob Storage
Maintainer&quot;" providerName="System.Data.SqlClient" />
    </connectionStrings>
    <RemoteBlobStorage>
        <Logging>
            <add key="ConsoleLog" value="0" />
        </Logging>
    </RemoteBlobStorage>
</configuration>
```

Figure 16

When complete, there should be a connection string for every content database in your farm. Copy this file so that it overwrites the old file in each client machine's maintainer directory.

Schedule RBS Maintainer tasks

Run a maintainer task on the web servers for each BLOB store for each directly connected web server.

For example, in the Hitachi Data Systems lab, the web server SP-WS01 has mapped BLOB stores for SPCDB01 and SPCDB02. Two maintainer tasks, one for each database and BLOB store mapping, are created on that server.

To create a basic task, do the following:

1. Type the **Name** of the title of the task and the **Description** for the task. When finished, click **Next**.

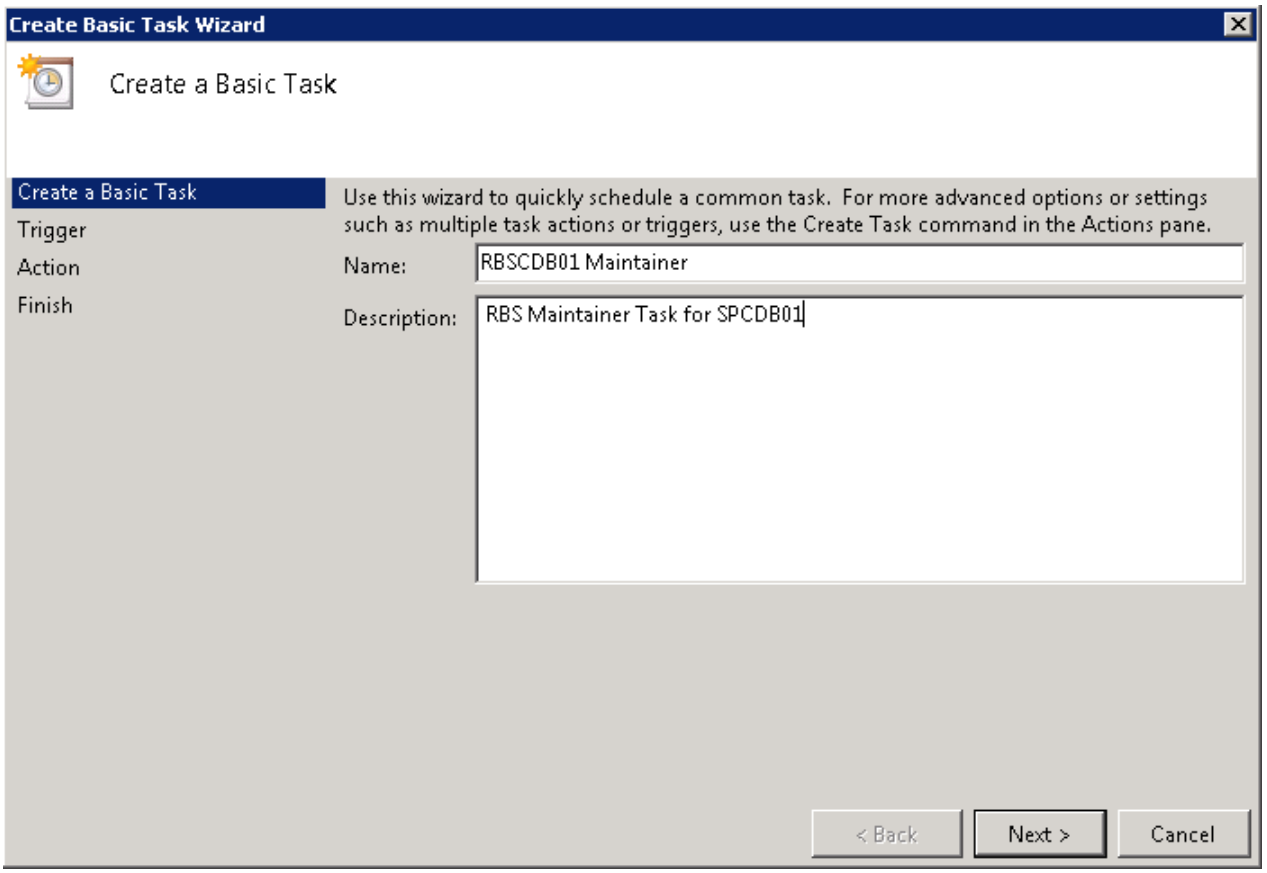

Figure 17

2. Select a daily trigger and click **Next**.
   - Run the maintainer tasks once a day, when server load is low..
3. Set the start time 2 hours apart from any other Maintainer task.. Click **Next**
   - Do not run maintainer tasks at the same time.
4. On the **Action** screen, verify that the **Start a program** button is selected and click **Next**.
5. In **Program/Script**, type the maintainer executable name:

```
C:\Program Files\Microsoft SQL Remote Blob Storage
10.50\Maintainer\Microsoft.Data.SqlRemoteBlobs.Maintainer.exe
```

6. In **Add arguments (optional)**, type in the following with the connection string name that was created in the **Microsoft.Data.SqlRemoteBlobs.Maintainer.exe.config**.

   ```
   ConnectionStringName <connection string name> -Operation
   GarbageCollection ConsistencyCheck ConsistencyCheckForStores -
   GarbageCollectionPhases rdo -ConsistencyCheckMode r -TimeLimit 120
   ```

7. In **Start in (optional)**, add the path to **Microsoft.Data.SqlRemoteBlobs.Maintainer.exe** and the click **Next**.

8. Click **Finish**.

Repeat this process to create a maintainer for each BLOB store on the web server where the BLOB store is mounted and shared.

In the Hitachi Data Systems lab, the RBS store volumes are distributed to the entire web front ends to balance the workload. When active, a sizable percentage of traffic to and from the BLOB stores is local to each web front end server. All of the RBS maintainer tasks are operating locally, further minimizing network traffic.

In the farm installed in the lab environment, half of the web front ends handle two BLOB stores and the associated maintainer tasks. The other half of the web front ends handle three BLOB stores and the associated maintainer tasks.

Once completing these are created, your SharePoint farm is ready to go live.

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