HITACHI
Inspire the Next

Microsoft SharePoint 2010 with Hitachi RBS Provider Adapter

Lab Validation Report

Mark Crowley

November 2011

Hitachi Data Systems

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

Lab Validation Report

This paper discusses testing of SAN-attached volumes on a Hitachi Virtual Storage Platform used for a SharePoint farm mapped using a UNC path on less expensive SATA volumes for BLOB storage.

Microsoft released new SharePoint 2010 content database sizing guidelines. Initially, the Microsoft guidelines called for the size of SharePoint 2010 content databases not to exceed 200GB. Currently, the Microsoft guidelines call for the size for SharePoint 2010 content databases not to exceed 4TB. See "SharePoint Server 2010 capacity management: Software boundaries and limits" on *TechNet*.

Following this guideline increases the possible maximum amount of unstructured data (uploaded data files) stored in a SharePoint SQL database. Even so, storing unstructured data in a SQL database remains an inefficient use of high performance storage in a typical SharePoint 2010 installation.

SharePoint stores unstructured data by default within the content databases on a SharePoint SQL server as BLOBs (**b**inary **l**arge **ob**jects). These BLOBs typically represent a large percentage of the stored data. To ensure maximum performance, the highest performance LUNs available are assigned to the administrator for BLOB storage.

RBS (remote BLOB storage) detaches BLOBs from the SQL database. With RBS storage, the database only stores the reference metadata for the BLOBs, with the BLOBs written to an alternate storage medium.

Using the RBS Provider for Hitachi Storage, you can assign BLOB storage to the following three SAN-attached storage mediums:

- CIFS UNC Path
- Hitachi Content Platform
- Hitachi Adaptable Modular Storage 2500

Note — This testing was done in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow recommended practice by conducting proof-of-concept testing for acceptable results before solution implementation in your production environment. This means to test applications in a non-production, isolated test environment that otherwise matches your production environment.

Test Environment Configuration

To demonstrate the use of RBS with SharePoint 2010, Hitachi Data Systems configured the following in its lab:

- A SharePoint 2010 farm with one SQL server
- Eight Web front ends and search server

Server Configuration

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

Table 1. Hardware Resources

Hardware	*Description*	*Quantity*	*Version*
Hitachi Virtual Storage Platform	▪ 4 × 8GB Fibre Channel ports ▪ 69GB cache memory ▪ 376 × 300GB 10k SAS ▪ 64 × 2TB 7.2K SATA	1	Microcode: 70-02-54-00/00 SVP: 70-02-53/00 RMI server: 06-02-01/00
Hitachi Compute Blade 2000 chassis	▪ 2 × 1/10Gb/sec Ethernet switch modules ▪ Used Blade 4, 5, 6, and 7	1	A0176-G-5666
Hitachi E55-A2 Server Blade	▪ 2 × Intel X5670 core processors (6 CPU cores per processor) ▪ 72GB RAM ▪ 2 × 300GB 10k SAS hard drives ▪ 1 × Mezzanine Gigabit Ethernet card ▪ 2 × Hitachi GVX-CC2N8G2X1 2-port PCIe HBA Fibre Channel ports	4	EFI v03-53
Brocade 5340 SAN Switch	▪ 80 8Gb/sec Fibre Channel ports	2	FOS 6.4.0a

The physical server configuration of the SharePoint farm is in Figure 1:

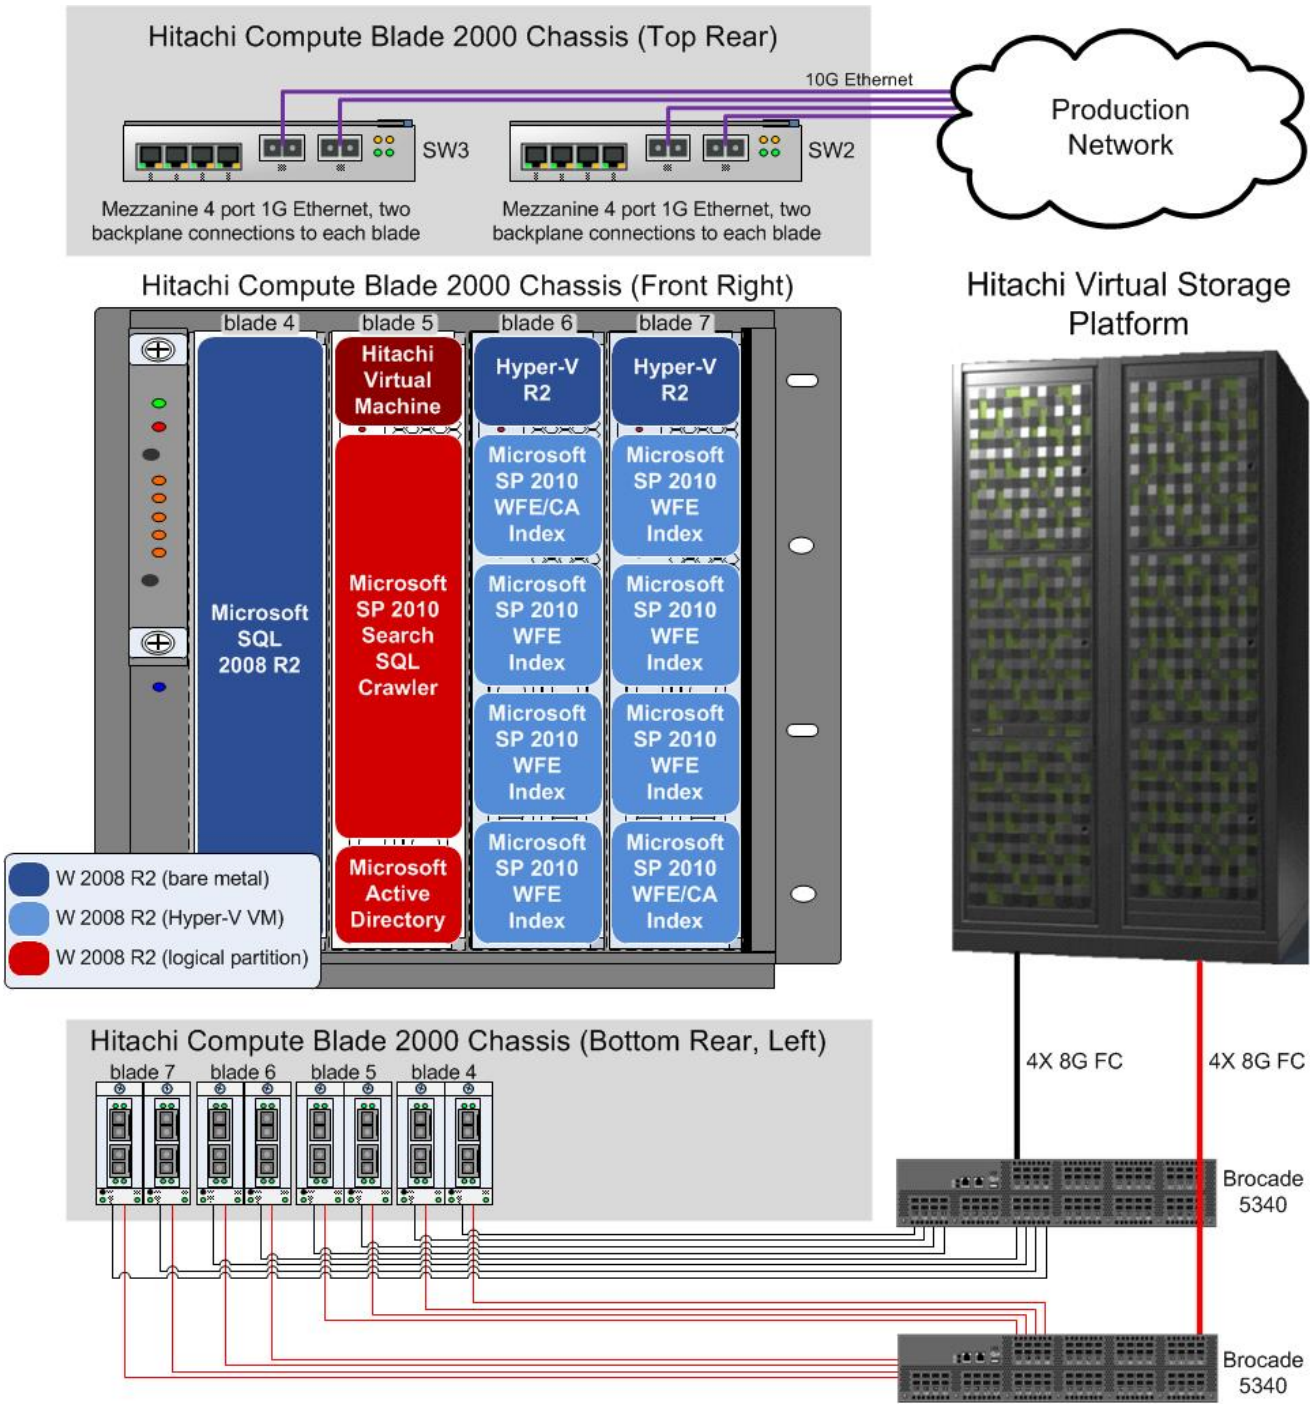
Hitachi Compute Blade 2000 Chassis (Top Rear)
10G Ethernet
Production Network
SW3
SW2
Mezzanine 4 port 1G Ethernet, two backplane connections to each blade
Mezzanine 4 port 1G Ethernet, two backplane connections to each blade
Hitachi Compute Blade 2000 Chassis (Front Right)
Hitachi Virtual Storage Platform
blade 4
blade 5
blade 6
blade 7
Microsoft SQL 2008 R2
Hitachi Virtual Machine
Microsoft SP 2010 Search SQL Crawler
Microsoft Active Directory
Hyper-V R2
Hyper-V R2
Microsoft SP 2010 WFE/CA Index
Microsoft SP 2010 WFE Index
Microsoft SP 2010 WFE Index
Microsoft SP 2010 WFE Index
Microsoft SP 2010 WFE Index
Microsoft SP 2010 WFE Index
Microsoft SP 2010 WFE Index
Microsoft SP 2010 WFE/CA Index
W 2008 R2 (bare metal)
W 2008 R2 (Hyper-V VM)
W 2008 R2 (logical partition)
Hitachi Compute Blade 2000 Chassis (Bottom Rear, Left)
blade 7
blade 6
blade 5
blade 4
4X 8G FC
4X 8G FC
Brocade 5340
Brocade 5340

Figure 1

Table 2. Blade Configuration

Blade	*Configuration*
4	Windows 2008 R2 Enterprise Edition (SP1) is installed directly on the blade server. This is for SharePoint use exclusively.
5	Hitachi Virtual Machine enabled on this server blade. Two logical partitions defined, as follows: ▪ LPAR 1 (8 CPUs/64GB RAM assigned —Windows 2008 R2 Enterprise Edition (SP1), Microsoft SQL Server 2008 R2, and Microsoft SharePoint 2010. SQL Server supports the search crawl database and the search property databases. ▪ LPAR 2 (4 CPUs/8GB RAM assigned)—Windows 2008 R2 Enterprise Edition (SP1), Microsoft Active Directory role (Windows 2008 R2 Forest/Domain) enabled.
6, 7	Microsoft Hyper-V enabled on each server blade. Each blade server has four virtual machines running Microsoft Windows 2008 R2 Enterprise Edition (SP1) and Microsoft SharePoint 2010. All virtual machines use pass-through disks (individual LUNs mapped directly from the host to the virtual machine) for boot and SharePoint install volumes.

Each virtual machine hosts two or three pass-through disks for the 20 BLOB store CIFS shares:

- SP-WS01 — RBSCDB01(\\SP-WS01\RBSCDB01), RBSCDB02(\\SP-WS01\RBSCDB02)
- SP-WS02 — RBSCDB03(\\SP-WS02\RBSCDB03), RBSCDB04(\\SP-WS02\RBSCDB04), RBSCDB05(\\SP-WS02\RBSCDB05)
- SP-WS03 — RBSCDB06(\\SP-WS03\RBSCDB06), RBSCDB07(\\SP-WS03\RBSCDB07)
- SP-WS04 — RBSCDB08(\\SP-WS04\RBSCDB08), RBSCDB09(\\SP-WS04\RBSCDB09), RBSCDB10(\\SP-WS04\RBSCDB10)
- SP-WS05 — RBSCDB11(\\SP-WS05\RBSCDB11), RBSCDB12(\\SP-WS05\RBSCDB12)
- SP-WS06 — RBSCDB13(\\SP-WS06\RBSCDB13), RBSCDB14(\\SP-WS06\RBSCDB14), RBSCDB15(\\SP-WS06\RBSCDB15)
- SP-WS07 — RBSCDB16(\\SP-WS07\RBSCDB16), RBSCDB17(\\SP-WS07\RBSCDB17)
- SP-WS08 — RBSCDB18(\\SP-WS08\RBSCDB18), RBSCDB19(\\SP-WS08\RBSCDB19), RBSCDB20(\\SP-WS08\RBSCDB20)

Note— A RBS maintainer task is defined on the machine where the BLOB stores are mapped, for each of the 20 RBS enabled databases.

Storage Architecture

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

Table 3. Hitachi Virtual Storage Platform Configuration

Hardware	*Description*	*Version*
Hitachi Virtual Storage Platform	▪ 4 × 8GB Fibre Channel ports ▪ 69GB cache memory ▪ 376 × 300GB 10k SAS ▪ 64 × 2TB 7.2K SATA	Microcode: 70-02-54-00/00 SVP: 70-02-53/00 RMI server: 06-02-01/00

Storage Pools

There are four dynamic provisioning storage pools, as described in Table 4:

Table 4. Dynamic Provisioning Pools

Dynamic Provisioning Pool Name (ID)	*RAID Group Config.*	*Drive Type*	*Number of RAID Groups*	*Pool Capacity (TB)*
HDT-MSFT-SP-DB(MT) (30) ▪ SQL database files (.mdf and .ndf) ▪ Dynamic provisioning pool using Hitachi Dynamic Tiering	RAID-5 (7D+1P)	300GB 10k RPM SAS ▪ Tier 1 for frequently accessed data	1	7.00
		2TB 7.2k RPM SATA ▪ Tier 2 for remaining data, using more cost-effective drives	2	
HDT-MSFT-SP-LOG (31) ▪ Dynamic provisioning pool used for the SQL database transaction log files (.ldf)	RAID-5 (7D+1P)	300GB 10k RPM SAS	1	1.87
HDT-MSFT-SP-BLOBSTORE(32) ▪ Dynamic provisioning pool used for the RBS BLOB stores that are presented to the SharePoint virtual machines.	RAID-5 (7D+1P)	2TB 7.2k RPM SATA	2	5.13
HDT-MSFT-SP-VMDATA (35) ▪ Dynamic provisioning pool used for the guest virtual machines boot and index volumes	RAID-5 (7D+1P)	300GB 10k RPM SAS	1	1.87

Table 5 through Table 13 provide details about the LDEVs used by the various hosts.

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

Table 5 lists the volumes provisioned for the configuration database.

Table 5. Configuration Database Volumes

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

Table 6 lists the volumes provisioned for the central administration database.

Table 6. Central Administration Database Volumes

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

Table 7 lists the volumes provisioned for the content databases.

Table 7. Content Databases Volumes

Dynamic Provisioning 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/4A/7A/8A
HDT-MSFT-SP-LOG (31)	31:02, 31:03, 31:04, 31:05, 31:06, 31:07, 31:08, 31:09, 31:0A, 31:0B, 31:0C, 31:0D, 31:0E, 31:0F, 31:10, 31:11, 31:12, 31:13, 31:14, 31:15	40	Logs 00-19	3A/4A/7A/8A

Table 8 lists the volumes provisioned for the search administration database.

Table 8. Search Administration Volumes

Dynamic Provisioning 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

Table 9. Crawl Database Volumes

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

Table 10. Search Property Database Volumes

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

Table 11. `tempDB` Volumes

Dynamic Provisioning 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 12 lists the volumes provisioned for the Hyper-V virtual machines and index volumes.

Table 12. Virtual Machine Volumes

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

Table 13 lists the volumes provisioned for the BLOB store volumes.

Table 13. BLOB Store Volumes

Dynamic Provisioning Pool Name (ID)	*LDEV*	*Size (GB)*	*Purpose*	*Storage Port*
HDT-MSFT-SP-BLOBSTORE (32)	30:00-30:09	200	SP-WS01—SP-WS04 BLOB stores (Blade6-HV)	3B/7B
HDT-MSFT-SP-BLOBSTORE (32)	30:0A-30:13	200	SP-WS05—SP-WS08 BLOB stores (Blade7-HV)	3B/7B

Test Methodology

All tests run were to verify RBS functionality. Upload, open, delete, and index operations were performed with various document types and sizes.

- A file explorer window for the BLOB store file system and the Hitachi RBS administration console are used to verify the presence and deletion of BLOBs from the BLOB store.
- A TSQL statement to change the BLOB file retention timeout was executed on the 20 content databases to cause the deletion of the BLOBs when the maintainer for that database executes, after file deletion on the 20 SharePoint sites.
- Testing was performed to verify RBS in the following functions:
- **Upload single document file—**Upload 15 files of various sizes and types to the SharePoint farm one at a time, verifying the creation of the BLOB and the verification of the BLOB. The document files were uploaded to the farm using the standard upload page. Documents were .docx, xlsx, pdf, and .mpg files.

- **Upload multiple document files —** Using the built in SharePoint functionality, upload 200 documents 20 files at a time using *Upload Multiple Documents* in SharePoint, verifying the alternate upload interface (WebDAV).
- **Overwrite single file** — Uploaded and overwrite an existing document with SharePoint versioning turned off, causing SharePoint to delete the original copy and write a new BLOB file. The RBS maintainer was scheduled in the task scheduler on the client machine,
- **Delete files —** Delete a group of files from SharePoint using the Datasheet view, observing what happens to the BLOB files in the BLOB stores. The RBS maintainer tasks for the BLOB stores were manually run on the client machines,
- **Run a search crawl —** Manually start the search system to read and index all BLOBs referenced by SharePoint and to look for errors.
- **Keyword search and open documents** — Search on the newly indexed documents to verify that the building of correct search indexes and that the correct documents open from search results page.

Analysis

Table 14 has test results and analysis.

Table 14. Test Results and Analysis

Test	*Results*	*Analysis*
Upload single document file	▪ No perceived change in performance or behavior from content database BLOB to remote BLOB storage. ▪ All files were readable in the content database and remote BLOB storage. ▪ The BLOB counter in the Hitachi administration console updated (increased) properly.	The BLOB file was created in the appropriate BLOB store and metadata was generated in the database. The Hitachi RBS Console counted the BLOB correctly.
Upload multiple document files	▪ No perceived change in performance or behavior from content database BLOB to remote BLOB storage. ▪ All files were readable in the content database and remote BLOB storage. ▪ The BLOB counter in the Hitachi administration console updated (increased) properly.	The BLOB files were created in the appropriate BLOB stores and metadata was generated in the database. The Hitachi RBS Console counted the BLOBs correctly.

Test	*Results*	*Analysis*
Overwrite Single File	▪ No perceived change in performance or behavior from content database BLOB to remote BLOB storage. ▪ RBS maintainer tasks were started manually on all of the client machines. ▪ The old document version in SharePoint deleted properly from the BLOB store. ▪ The BLOB counter in the Hitachi administration console updated (decreased) properly.	The new BLOB was written into the BLOB store, and the old version was removed. The RBS maintainer tasks were scheduled to run once an hour, but were started manually to speed up test execution.
Delete Files	▪ No perceived change in performance or behavior from content database BLOB to remote BLOB storage. ▪ RBS maintainer tasks were started manually on all of the client machines. ▪ The deleted document in SharePoint deleted properly from the BLOB store. ▪ The BLOB counter in the Hitachi administration console updated (decreased) properly.	No errors recorded. Documents were removed from the BLOB stores. The RBS maintainer tasks were scheduled to run once an hour, but were started manually to speed up test execution.
Run a Search Crawl	▪ Performance was completely normal. ▪ All documents on the site were properly indexed.	No errors reading BLOBs recorded during search crawl
Keyword Search and Open Documents	▪ Search returned normal results ▪ Documents referenced by search returns opened correctly,	Open BLOBs referenced by search returns was successful.

The Hitachi RBS Provider, coupled with a single Microsoft SharePoint web front end and Microsoft SQL Server performed properly. The performance difference was indistinguishable compared to a database without RBS enabled.

Appendix — Remote BLOB Storage Overview

This briefly describes how BLOB storage works. From more information, see *Overview of RBS (SharePoint Server 2010)* on Microsoft *TechNet*.

RBS Architecture

The RBS Components and their relationships are shown in Figure 2.

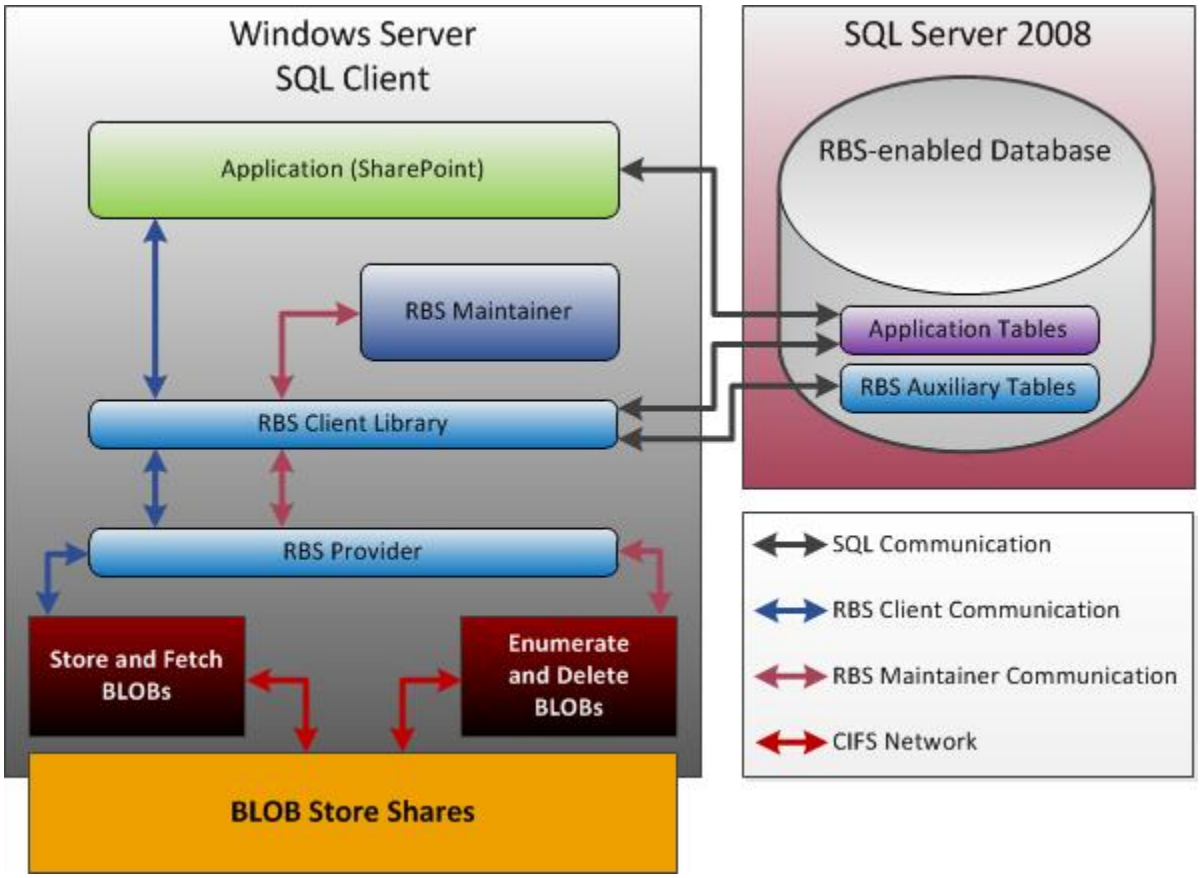

Figure 2

SQL Server

RBS is a feature new with Microsoft SQL 2008. It stores BLOB data outside of a database. The Microsoft RBS feature pack install creates RBS auxiliary tables and stored procedures in a targeted database. The installed client components maintain RBS columns in application tables to write BLOB references to the databases.

Client Application

The client application interacts with the RBS client library provided by Microsoft to access the BLOB store. To create a BLOB, an application invokes the RBS client library and passes the database connection object and the data to be stored in the BLOB to the client library

RBS Client Library

Microsoft provides the RBS client libraries as part of Microsoft SQL Server 2008 feature pack. RBS client library does the following to perform the client application request:

- Finds the provider information from the SQL Server database
- Loads the provider library
- Invokes the respective API

When creating a BLOB, it stores the BLOB reference in RBS auxiliary tables and returns the BLOB-ID to the client application. The client application stores the BLOB-ID in RBS column in application tables.

RBS Provider

It is a managed library which implements the BlobStore abstract class provided by RBS libraries from Microsoft. It interacts with the BLOB store for storing and retrieving BLOBs. Hitachi developed a custom RBS provider library to interface with BlobStore abstract for storing and retrieving BLOBs.

RBS Maintainer

RBS Maintainer is a command line application which is provided by Microsoft to perform the BLOB Store cleanup process (garbage collection). Add Maintainer to Windows Scheduler to run at configured times. Install the Hitachi RBS Provider on each RBS Maintainer machine.

RBS Maintainer does garbage collection in 3 phases:

- **Reference Scan (RS)** — Find BLOBs in the application tables that are no longer referenced by client application. This uses the list of registered RBS columns. The BLOBs that are no longer referenced by the application are marked to be deleted.
- **Delete Propagation (DP)** — Delete BLOBs marked for deletion from the BLOB store the number of days set in configuration. There is a gap between when a BLOB is marked for deletion and when a BLOB actually is deleted. Configure the number of days for gap using <garbage_collection_time_window>.The default is 30 days.
- **Orphan Cleanup (OC)** — Identify all BLOBs in the BLOB store and determine the list of BLOB files that are present in the BLOB store but are not referenced in the RBS tables within the RBS-enabled database. These BLOBs, called orphans, occur from aborted transactions, application misbehavior, or other failures. The maintainer deletes the orphan BLOBs that were created before the garbage collection time window from the BLOB store.

Consistency Check, another important function of Maintainer, verifies the following:

- Auxiliary tables are in consistent state and each RBS table column has a valid index.
- Registered-application RBS tables columns exist in the database, have been RBS enabled, and have valid indexes.
- No BLOBs are marked as in use and deleted at the same time.
- All BLOBs that are referenced in the application tables are present in the RBS tables.

Any discovered problems will be logged and the RBS Maintainer will attempt to fix them by creating missing index entries, unregistering missing columns, or marking in-use BLOBs as not deleted.

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