HITACHI
Inspire the Next

Deploying SQL Server 2008 R2 with Hyper-V on the Hitachi Virtual Storage Platform

Implementation Guide

By Rick Andersen

November 2010

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 SQL Server 2008 R2 with Hyper-V on the Hitachi Virtual Storage Platform

Implementation Guide

IT administrators need database solutions that provide data protection and simple management in environmentally friendly data centers. Using Microsoft SQL Server 2008 R2 with SAN storage like the Hitachi Virtual Storage Platform accomplishes those business-critical objectives. This implementation guide leverages the features of SQL Server 2008 R2 and the advantages of the Virtual Storage Platform to lower administration costs, improve system efficiency and enable virtualization.

SQL Server 2008 R2 can be deployed for many types and numbers of users in a wide variety of infrastructure topologies. By implementing SQL Server 2008 R2 with the Virtual Storage Platform, you can easily scale from a small environment to a very large environment. .

The Hitachi Virtual Storage Platform can help you leverage your information, which is the new currency in today's data-driven economy. Information exists in many forms and must be protected and readily accessible to ensure business survival and success. The Virtual Storage Platform maximizes cost efficiency and return on investment by creating an agile storage infrastructure that reduces costs and increases performance, availability, scalability and reliability.

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

In addition, Hitachi Dynamic Provisioning software also allows you to combine different workloads in a single storage frame for greater flexibility. Hitachi Dynamic Provisioning software's thin provisioning and wide striping functionalities provide virtual storage capacity to reduce costs and simplify administration.

This implementation guide focuses on deploying SQL Server 2008 R2. It is based on the reference architecture described in the *Deploying SQL Server 2008 R2 with Hyper-V on the Hitachi Virtual Storage Platform Reference Architecture Guide* white paper. It is intended for use by IT administrators responsible for SQL and storage. It assumes some familiarity with Hitachi Storage Navigator software, Microsoft Windows Server 2008 R2 and SQL Server 2008.

Solution Overview

This white paper describes how to deploy a SQL Server 2008 R2 environment that supports 75,000 users on a Virtual Storage Platform storage system. For ease of management and scalability and to provide predictable performance, this solution uses a building block approach. A single virtual machine (VM) running Windows Server 2008 R2 and SQL Server 2008 R2 with underlying storage from the Virtual Storage Platform make up a building block. The Virtual Storage Platform takes advantages of Hyper-V features such as live migration via storage replication technologies to establish a reliable and highly available virtualized SQL Server solution. This solution scales from a 75,000-user deployment with a 900GB database to a 150,000-user deployment with a second 900GB database by implementing a second virtual machine and SQL Server instance.

The building block for the 75,000-user account database consists of two 450GB LUNs for the SQL database, a 90GB LUN for tempdb and a 225GB LUN for the SQL transaction log. This building block was replicated along with adding a second VM and SQL Server instance to support 150,000 users on the Virtual Storage Platform. All LUNs were provisioned from Dynamic Provisioning pools on the Hitachi Virtual Storage Platform; note that passthrough disks were used for the database, `tempdb` and log LUNs.

Figure 1 shows the high level storage design of the SQL Server 2008 R2 reference architecture on the Virtual Storage Platform.

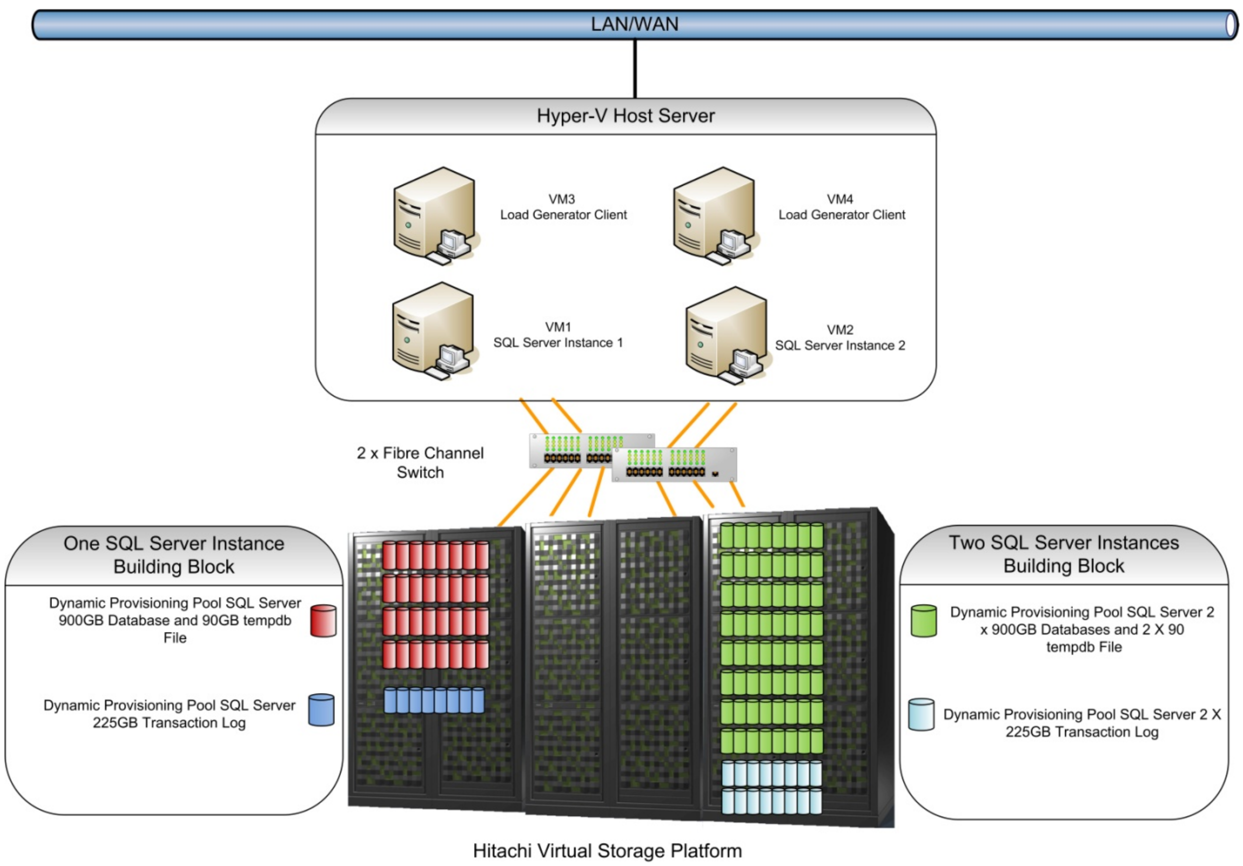

Figure 1

Hitachi Virtual Storage Platform

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

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

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

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

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

The switch matrix architecture of the Virtual Storage Platform makes all of this possible. It connects the basic components, front-end directors, back-end directors, global cache modules and virtual storage directors. You can add redundant pairs of directors and cache modules as required without disruption to connected host servers. All these resources are tightly coupled through a global cache that creates a common pool of storage resources. These resources can include external storage that is connected through front-end director initiator ports.

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

Hitachi Dynamic Provisioning Software

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

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

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

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

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

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

Microsoft Windows Hyper-V

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

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

Microsoft SQL Server 2008 R2 Software

Microsoft SQL Server 2008 R2 provides a scalable, high performance database engine for mission-critical applications that require the highest levels of availability and security. SQL Server 2008 R2 also provides enhanced enterprise-class manageability for large OLTP deployments like the one that is described in this implementation guide. Together with the Hitachi Virtual Storage Platform, SQL Server 2008 R2 provides a scalable, high-performance database engine for any midrange to enterprise level application.

Tested Solution Components

The following sections detail the hardware and software components that were deployed in the Hitachi Data Systems lab.

Hardware Components

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

Table 1. Hardware Components

Hardware	*Detail Description*	*Version*	*Quantity*
Hitachi Virtual Storage Platform storage system	Single frame 6x 8GB Fibre Channel ports used 128GB cache memory 72 x 300GB, 10K RPM, SAS disks	70-00-50-00/00	1
Brocade 5300 switch	SAN switch with 8Gb Fibre Channel ports	FOS 6.4.0b	2
Dell PowerEdge R905 server	4 x Quad-Core AMD Opteron processor 1.9 GHz, 128GB RAM. Equipped with 2 x Emulex LPe12002 8GB HBAs	BIOS firmware 4.2.1	1

For this solution, Hitachi Data Systems used Hitachi Dynamic Provisioning software to provision storage. Table 2 lists storage configuration details for the Dynamic Provisioning pools created for the SQL Server building block supporting 75,000 users. RAID groups used to create the Dynamic Provisioning pools were RAID-1+0 (2D+2D) and RAID-5 (3D+1P). This building block was replicated to support the 150,000 user environment.

Table 2. Dynamic Provisioning Pool for Single Virtual Machine Supporting 75,000 Users

Pool	*Number of Drives*	*Drive Capacity*	*Number of LUNs per Pool*	*Role*
01	32	300GB SAS 10K RPM	3	SQL Server database and `tempdb` files
02	8	300GB SAS 15K RPM	1	SQL Server transaction logs
03	4	2000GB SATA 7200 RPM	1	Guest OS VHD in a shared pool with other guest OS VHDs

Table 3 lists storage configuration details for LDEVs created for the SQL Server instance that supports a single virtual machine supporting 75,000 users. All LDEVs for the SQL Server instance were mapped to storage ports 1C, 2C, 5C and 6C. The LDEV for the guest OS VHD was mapped to storage ports 3B and 4B. This was replicated to support a 150,000 user environment.

Table 3. LDEV Storage Configuration for Single Virtual Machine Supporting 75,000 Users

Pool	*LDEV*	*LUN*	*Size (GB)*
03	00:02:00	03	150
01	00:20:10	04	450
01	00:20:11	05	450
01	00:20:12	06	90
02	00:20:02	07	225

Table 4 lists the LUN configuration for a single building block supporting 75,000 users. This same LUN configuration is replicated on a second SQL Server guest machine to support the 150,000 user architecture. Note: The LUNs for the SQL Server operating system are presented and mounted on the Hyper-V host only.

Table 4. LUN Configuration

LUN	*Size*	*Description*	*Pool ID*
3	150GB	OS LUN for SQL Server VM Operating System	03
4	450GB	Database LUN1 for SQL Server	01
5	450GB	Database LUN 2 for SQL Server	01
6	90GB	`tempdb` for SQL Server	01
7	225GB	Transaction logs for SQL Server	02

Storage Area Network Components

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

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

Figure 2 shows the Fibre Channel SAN architecture for the SQL Server 2008 R2 implementation described in this solution.

Figure 2

Software Components

Table 5 lists the detailed information about the software components used.

Table 5. Software Components

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

Table 6 lists the detailed information about the virtual machines used.

Table 6. VM Servers on Physical Host

VM Server Name	*Role*	*Number of vCPUs*	*Virtual Memory (GB)*
H-DC1	Active Directory domain controller	1	2
H-SQL-01	SQL Server 2008 R2 instance 1	4	32
H-SQL-02	SQL Server 2008 R2 instance 2	4	32

Solution Implementation

To deploy this SQL Server 2008 R2 solution, follow these high-level steps:

1. Configure fabric switch zones.
2. Configure the Virtual Storage Platform.
3. Install SQL Server 2008 R2.
4. Configure SQL Server 2008 R2.

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

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

- Hitachi Virtual Storage Platform online help
- Hitachi Virtual Storage Platform Dynamic Provisioning Software user's guide
- Microsoft TechNet library for SQL Server 2008 R2.

Configure Fabric Switch Zones

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

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

Configure the Virtual Storage Platform

The following sections describe how to configure the Virtual Storage Platform for this solution.

Creating a Dynamic Provisioning Pool

You must create at least two pools on the Virtual Storage Platform, one for the SQL Server databases and tempdb and one for the SQL Server transaction logs.

This procedure assumes that zoning within your Fibre Channel fabric is complete. To create a Dynamic Provisioning pool using Hitachi Storage Navigator software, follow these steps:

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

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

2. From the **Pool Type** drop-down menu, choose **Dynamic Provisioning**.
3. Select the **Disable** radio button for the **Multi-Tier Pool** option.
4. (Optional) Choose a menu item from the **Drive Type/RPM** drop-down menu and from the **RAID Level** drop-down menu.

   These options allow you to filter the available pool volumes.

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

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

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

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

7. Assign a prefix for the pool name in the **Prefix** field.
8. (Optional) Assign an initial number for the pool name in the **Initial Number** field.
9. Expand the **Options** pane.
10. Assign a pool id in the **Pool ID** field.
11. Assign a subscription limit in the **Subscription Limit** field.

    This is the percentage of oversubscription for this pool that you allow in your environment.

12. Choose a value from the **User-Defined Threshold** drop-down menu and click the **Add** button.

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

    The **Selected Pools** pane is populated.

13. Click the **Finish** button.

    The **Create Pools Dialogue** box displays.

14. Click the **Apply** button.

Create LDEVs within a Dynamic Provisioning Pool

This procedure assumes that Dynamic Pool creation in your environment is complete. To create one or more LDEVs within a Dynamic Provisioning pool using Hitachi Storage Navigator software, follow these steps:

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

   The **Create LDEVs** window displays.

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

3. From the **Emulation Type** drop down menu, select **OPEN-V**.

4. In the **Multi-Tier Pool** field, select the **Disable** radio button.

5. (Optional) Choose a menu item from the **Drive Type/RPM** drop-down menu and from the **RAID Level** drop-down menu.

   These options allow you to filter the available pools.

6. Click the **Select Pool** button.

   The **Select Pool** window displays.

7. Highlight a pool in the **Available Pools** pane and click **OK**.

   The **Create LDEVs** window displays with the **Selected Pool Name** and the **Selected Pool Capacity** fields populated.

8. Enter a capacity amount in the **LDEV Capacity** field and choose a unit of measure from the drop-down menu.

9. Enter the number of LDEVs of that size to be created in the **Number of LDEVs** field.

10. In the **LDEV Name** pane, assign a prefix in the **Prefix** field and assign an initial number in the **Initial Number** field.

11. Expand the **Options** pane.

12. Review the value in the **LDKC** field.

    Modify the LDKC value if the default of **`00`** is not appropriate. This is most often the case if the storage will be configured with more than one LDKC.

13. Choose a value from the **CU** drop-down menu.

14. Choose a value from the **DEV** drop-down menu.

15. (Optional) Choose a value from **Interval** drop-down menu.

    Leave this value at the default of **0** for sequential numbering of LDEVs. If you want a different numbering sequence, choose a different value.

16. Review the default values in the **Initial SSID** field, the **CLPR** field and **Processor Blade** field.

    In most situations, use the default values. Change them only if your environment requires different values.

17. Click the **Add** button.

    The **Selected LDEVs** pane is populated.

18. Click the **Finish** button.

    The Create **LDEVs** window displays.

19. Click the **Apply** button.

Creating a Host Group

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

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

   The Create Host Groups window displays.

2. Assign a name in the **Host Group Name** field.

3. From the **Host Mode** drop-down menu, choose 2C[Windows Extended].

4. In the **Available Hosts** pane, highlight one or more hosts.

5. In the **Available Ports** pane, highlight one or more ports.

6. Click the **Add** button.

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

7. Click the **Finish** button.

   The **Create Host Groups** window displays.

8. Click the **Apply** button.

Mapping an LDEV

LDEV mapping is the process through which storage administrators make LDEVs available to hosts and servers so that data can be stored and accessed by the host and its applications.

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

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

   The **Add LUN Paths** window displays.

2. In the **Available LDEVs** pane, highlight one or more LDEVs.

3. Click the **Add** button.

   The **Selected LDEVs** pane is populated.

4. Click the **Next** button.

   The **Add LUN Paths** window displays.

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

6. Click **Add**.

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

7. Click **Next**.

   The **Add LUN Paths** window displays.

8. Click the **Finish** button.

   The **Add LUN Paths** window displays.

9. Click the **Apply** button.

Install SQL Server 2008 R2

To install SQL Server 2008 R2, follow Microsoft's instructions in TechNet's [SQL Server 2008 Library] under the SQL Server 2008 R2 section.

Configure SQL Server 2008 R2

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

Databases and Transaction Log Files

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

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

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

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

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

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

`tempdb` Files

The default location for tempdb files in SQL Server is on the local server `C:` drive. You can modify the existing `tempdb` file location to point to the LUN provisioned on the Virtual Storage Platform. This can be done by executing the following SQL statements:

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

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

```
USE tempdb
GO
ALTER DATABASE tempdb
MODIFY FILE (NAME='tempdev', FILENAME= 'T:\SQL\DATA\datatempdb1.mdf',
SIZE=8192 KB, FILEGROWTH=42MB)
```

Note that the files are now located on a LUN that is assigned to the Hyper-V server for each SQL virtual machine's `tempdb`.

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