TECHNICAL PAPER

HITACHI
Inspire the Next

Deploy Hitachi Unified Compute Platform Select for Microsoft® SQL Server® 2008 in a Fast Track Data Warehouse Solution

Implementation Guide

By Nathan Tran

November 28, 2012

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

Deploy Hitachi Unified Compute Platform Select for Microsoft® SQL Server® 2008 in a Fast Track Data Warehouse Solution

Implementation Guide

This implementation guide focuses on deploying Hitachi Unified Compute Platform Select for Microsoft SQL Server 2008 in a Fast Track Data Warehouse solution. It is based on the reference architecture found in *Deploy Hitachi Unified Compute Platform Select for Microsoft® SQL Server® 2008 in a Fast Track Data Warehouse Reference Architecture Guide*.

Microsoft SQL Server 2008 R2 is a comprehensive database server and information platform to help organizations derive the most value from business data. This solution reduces the complexity to implement a Fast Track Data Warehousing solution by using servers and storage from Hitachi that work together seamlessly while providing a high level of redundancy.

This implementation guide is intended for you if you are an administrator deploying Microsoft SQL Server 2008 R2 for your organization on Hitachi Unified Storage 150 and Hitachi Compute Rack 220. Familiarity with the following is necessary to benefit from this implementation guide:

- Hitachi Unified Storage 100 family
- Hitachi Storage Navigator Modular 2
- Microsoft Windows Server 2008 R2
- Microsoft SQL Server 2008 R2

Solution Components

The following are the components used in this solution:

- **Hitachi Compute Rack 220**—2U rack-mountable server
- **Hitachi Unified Storage 150**—High performance and scalable storage system
- **Emulex LightPulse 8 Gb/sec Dual Port Fibre Channel HBA**—High performance connectivity to the storage network
- **Microsoft Windows Server 2008 R2**—Multi-purpose server software designed to increase the reliability and flexibility of your infrastructure
- **Microsoft SQL Server 2008 R2**—Enterprise-ready database technologies and tools that help business derive the most value from information at a high level of performance, scalability and availability.

Figure 1 illustrates a single node Microsoft SQL Server 2008 R2 environment using Hitachi Compute Rack 220 and Hitachi Unified Storage 150.

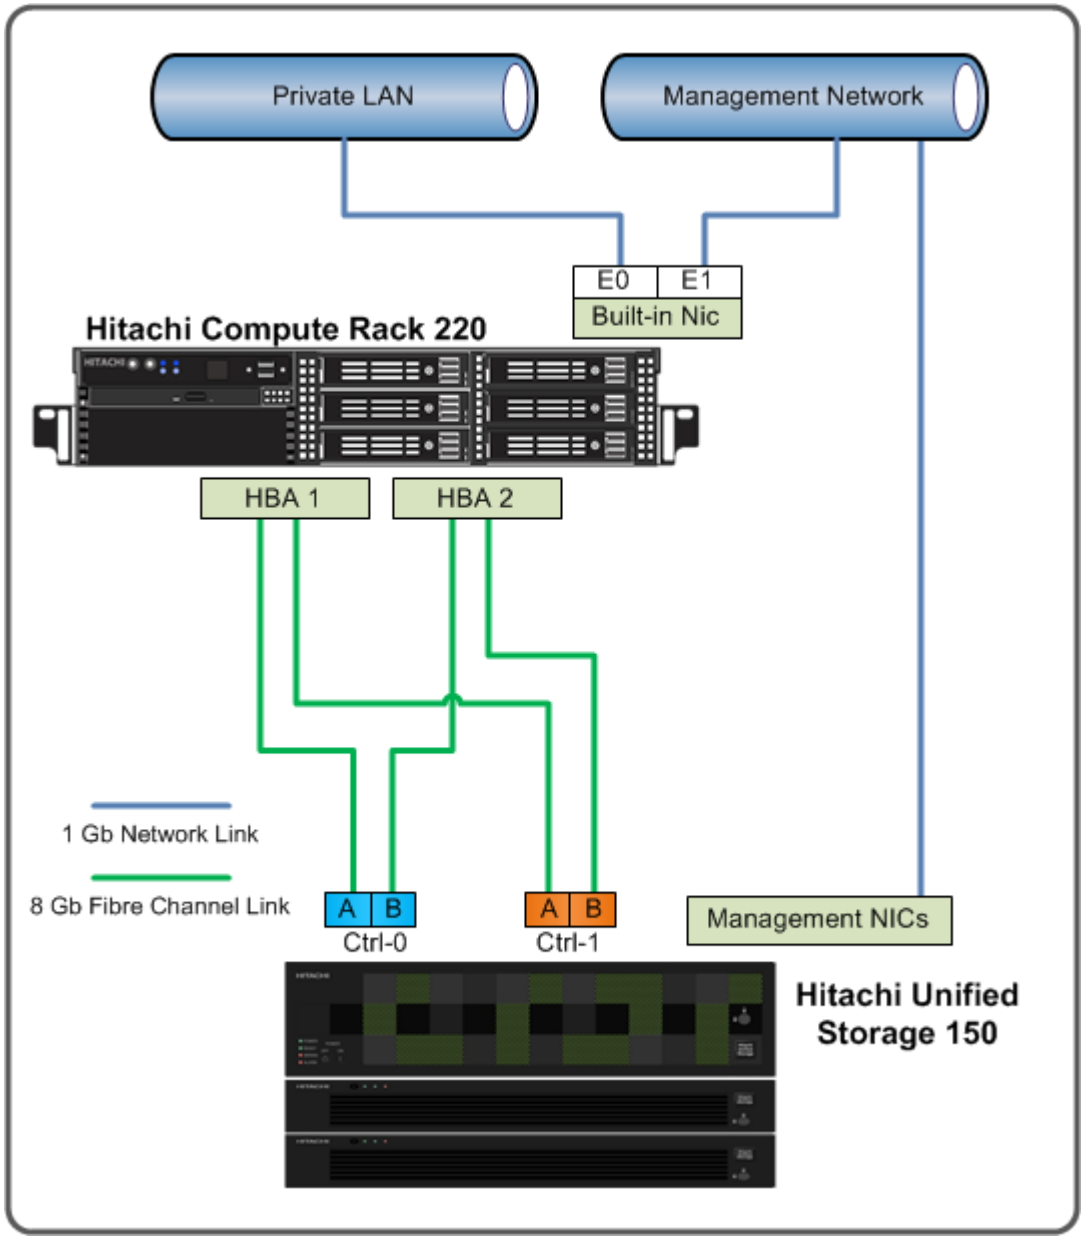

Figure 1

Hardware Components

Table 1 lists the hardware used for this solution.

Table 1. Hardware Components

Component	Description	Firmware	Quantity
Server	Hitachi Compute Rack 220 ■ Form Factor 2U (rack mountable) ■ 2 × Quad-Core Intel Xeon 2.4 GHz Processor E5620 ■ 96 GB RAM (12 × 8 GB DDR3 RDIMM) ■ 2 ×1000BASE-T/100BASE-TX/10BASE-T ports	7TTSHE-F9	1
Storage System	Hitachi Unified Storage 150 ■ 48 × 900 GB 10k RPM SAS drives (19.2 TB initial capacity) ■ 2 × SFF Dense Disk Expansion Tray	08B5/A-Z	1
Host Bus Adapter	Emulex LightPulse LPe12002-M8 ■ 8 Gb/sec dual port ■ PCI Express 2.0	2.50.007-3	2

Hitachi Unified Storage 150

Hitachi Unified Storage is a midrange storage platform for all data. It helps businesses meet their service level agreements for availability, performance, and data protection.

The performance provided by Hitachi Unified Storage is reliable, scalable, and available for block and file data. Unified Storage is simple to manage, optimized for critical business applications, and efficient.

Using Unified Storage requires a smaller capital investment. Deploy this storage, which grows to meet expanding requirements and service level agreements, for critical business applications. Simplify your operations with integrated set-up and management for a quicker time to value.

Unified Storage enables extensive cost savings through file and block consolidation. Build a cloud infrastructure at your own pace to deliver your services.

Hitachi Unified Storage 150 provides reliable, flexible, scalable, and cost-effective modular storage. Its symmetric active-active controllers provide input-output load balancing that is integrated, automated, and hardware-based.

Both controllers in Unified Storage 150 dynamically and automatically assign the access paths from the controller to a logical unit (LU). All LUs are accessible, regardless of the physical port or the server that requests access.

Hitachi Compute Rack 220

Hitachi Compute Rack 220 is an Intel Xeon processor-based midrange rack mountable server platform, providing advanced systems management and redundancy options. It is data center friendly, with a 2U footprint, while delivering the performance that is required to meet enterprise level challenges.

The benefits of Hitachi Compute Rack 220 are the following:

- Web-based management interface
- RAID level configuration, with up to six 3.5 inch internal drives
- Eco-friendly power-saving capabilities
- 2 socket Intel based server
- Configuration flexibility to meet business needs
- Dense 2U rack mountable design

This reference architecture uses Hitachi Compute Rack 220 for the following reasons:

- High processing power in a dense 2U design
- Meets the requirements of the Microsoft SQL Fast Track program for processor, RAM, and network capability

Software Components

Table 2 describes the software used for this solution.

Table 2. Software Components

Software Components	*Version*
Hitachi Storage Navigator Modular 2	Microcode Dependent
Hitachi Dynamic Link Manager Advanced	6.6.0-00
Microsoft Windows Server	2008 R2 SP1, Enterprise edition (64-Bit)
Microsoft SQL Server	2008 R2 SP1, Enterprise edition (64-Bit)

Hitachi Storage Navigator Modular 2

Hitachi Storage Navigator Modular 2 provides essential management and optimization of storage system functions. Using Java agents, Storage Navigator Modular 2 runs on most browsers. A command line interface is available.

Use Storage Navigator Modular 2 for the following:

- RAID-level configurations
- LUN creation and expansion
- Online microcode updates and other system maintenance functions
- Performance metrics

Hitachi Dynamic Link Manager Advanced

Hitachi Dynamic Link Manager Advanced combines all the capabilities of Hitachi Dynamic Link Manager and Hitachi Global Link Manager into a comprehensive multipathing solution. It includes capabilities such as the following:

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

The configuration of each data and log LUN has a primary active link and a secondary failover link. There is an even distribution of all LUNs among the four Fibre Channel paths. In case of a link failure, the affected LUN switches to another active link to maintain connectivity.

Microsoft recommends this practice for a Fast Track Data Warehouse solution.

Microsoft Windows Server 2008 R2

Microsoft Windows 2008 Server R2 is a multi-purpose server designed to increase the reliability and flexibility of your server or private cloud infrastructure.

Additional highlights of Microsoft Windows Server 2008 R2 include the following:

- Develop, deliver, and manage rich user experiences and applications.
- Provide a highly secure network infrastructure.
- Increase technological efficiency and value within your organization.

For more information, see “Product Information” on the Microsoft website.

Microsoft SQL Server 2008 R2

Microsoft SQL Server 2008 R2 is a complete set of enterprise-ready database technologies and tools to help your business derive the most value from information. SQL Server provides high levels of performance, availability, and security while employing more productive management and application development tools.

Additional highlights of Microsoft SQL Server 2008 include the following:

- A complete business intelligence platform that connects users to the right information at the right time. This improves business decisions through familiar tools, such as Microsoft Excel and Microsoft SharePoint Server.
- You receive high levels of performance, scalability, availability, and security for mission-critical applications.
- With Microsoft Visual Studio, the Microsoft .NET Framework, and Microsoft SQL Server, developers use integrated development tools to build rich, intuitive, and connected applications quickly.

For more information about the features of SQL Server 2008 R2, see the What's New page of SQL Server 2008 R2 Books Online or "Product Information" on the Microsoft web site.

Solution Implementation

Deploying this solution requires the following procedures.

1. “Configure Hitachi Unified Storage 150” on page 8
2. “Configure Hitachi Compute Rack 220” on page 15
3. “Configure Storage Area Network (SAN)” on page 19
4. “Install Software ” on page 20
5. “Configure Microsoft SQL Server 2008 R2” on page 21

Your implementation checklist may vary based on your environment and business requirements.

Configure Hitachi Unified Storage 150

Configure the storage using the following procedures.

Configure Fibre Channel Port Settings

To configure the Fibre Channel port settings, do the following for Port 0A, Port 1A, Port 0B, and Port 1B.

1. Open Hitachi Storage Navigator Modular 2.

   (1) Open a web browser and type the following in the address bar:

   `http://127.0.0.1:23015/StorageNavigatorModular/Login`

   (2) Log on using these credentials.

   User ID: `system`

   Password: `manager`

2. Open the Fibre Channel port settings.

   (1) Click the **Array Name** link to open the storage system.

   (2) Expand the **Settings** heading and click the **FC Settings** link.

   (3) Click **FC Port <port name>**.

   The Fibre Channel port properties display with an option to edit the Fibre Channel port.

3. Edit the Fibre Channel port settings.

   (1) Click **Edit FC Port**.

   (2) Click **8Gbps** from the **Transfer Rate** menu.

   (3) Click **Loop** from the **Topology** menu.

   (4) Click **OK**.

   A message displays indicating that the change interrupts I/O between any hosts that are connected to the port at this time.

   (5) Click **Confirm** and wait a few seconds for the change to take place.

After establishing the connection between the storage system and the host, the **FC Settings** window shows all configured ports with the status **LinkUp(Private Loop)**.

Create RAID Groups

This solution uses RAID-5 (4D+1P) RAID groups for OS boot and RAID-10 (2D+2D) for all SQL RAID groups. Detailed RAID group configuration is in Table 3.

Table 3. RAID Group Configuration

RAID Group	*RAID Type*	*Number of Drives*	*Usable Pool Capacity (TB)*	*Purpose*
0	RAID- 5 (4D+1P)	5	3.2	SAN OS Boot
1	RAID-10 (2D+2D)	4	1.6	Database, tempdb
2	RAID-10 (2D+2D)	4	1.6	Database, tempdb
3	RAID-10 (2D+2D)	4	1.6	Database, tempdb
4	RAID-10 (2D+2D)	4	1.6	Database, tempdb
5	RAID-10 (2D+2D)	4	1.6	Database, tempdb
6	RAID-10 (2D+2D)	4	1.6	Database, tempdb
7	RAID-10 (2D+2D)	4	1.6	Database, tempdb
8	RAID-10 (2D+2D)	4	1.6	Database, tempdb
9	RAID-10 (2D+2D)	4	1.6	Transaction log
10	RAID-10 (2D+2D)	4	1.6	Transaction log

To create each RAID group using the RAID group configuration information in Table 3, follow these steps, starting with RAID group 0.

1. Open Hitachi Storage Navigator Modular 2.

   (1) Open a web browser and type the following in the address bar:

   `http://127.0.0.1:23015/StorageNavigatorModular/Login`

   (2) Log on using these credentials.

   User ID: `system`

   Password: `manager`

2. Open the volumes settings.

   (1) Click the **Array Name** link to open the storage system.

   (2) Expand the **Groups** heading and click the **Volumes** link.

   The right pane has three tabs: **Volumes**, **RAID Groups**, and **DP Pools**.

3. Create a RAID Group.

   (1) Click the **RAID Groups** tab and then click **Create RG**.

   The **Create RAID Group** window displays.

   (2) Click the following from each menu, as shown in Figure 2.

   - *RAID5* from the **RAID Level** menu
   - *4D+1P* from the **Combination** menu

Select the appropriate settings by referencing Table 3 on page 9 when creating other RAID groups.

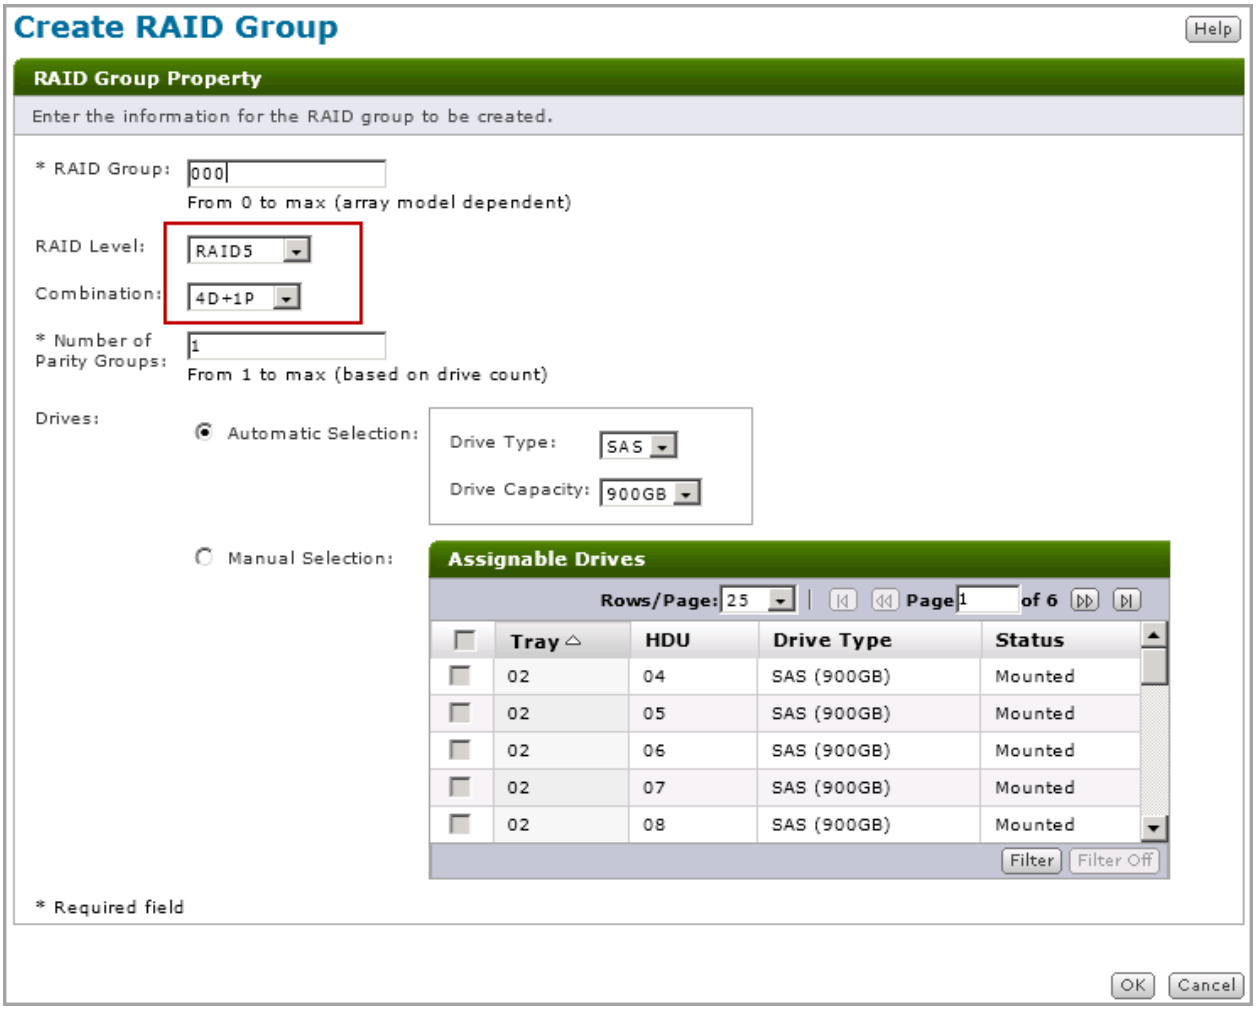

Figure 2

The **Number of drives** changes automatically, based on what you click for the RAID level and combination.

(3) Click an option for **Drives**.

- **Automatic Selection**—Click this option to have Hitachi Storage Navigator Modular 2 select the next available drives shown in the **Assignable Drives** table.
- **Manual Selection**—Click this option if you have different types of drives installed in the storage system. These could be SATA and SAS drives, or drives of different capacities. Then select the correct drive type in the **Assignable Drives** table so that Hitachi Storage Navigator Modular 2 selects the correct type of hard drive.

Hitachi Data Systems recommends clicking **Automatic Selection** when all your drives are the same type.

(4) Click the **Advanced** tab to modify any of the settings based on your environment requirements, and then click **OK**.

A message displays indicating successful creation of the RAID group.

4. Click **Close**.

The RAID group immediately starts the formatting process in the background.

Use this process to create each RAID group in Table 3 on page 9.

Create Volumes

Create volumes from the RAID groups for presenting to the host for the following purposes:

- SAN OS boot
- Stage Volume
- SQL Server database and tempdb
  - DB 1
  - DB 2
  - DB 3
  - DB 4
  - DB 5
  - DB 6
  - DB 7
  - DB 8

- SQL Server Transaction Log volumes
  - LOG 1
  - LOG 2

Table 4 has the volume configuration for this solution.

Table 4. Volume Configuration

RAID Group	*Volume*	*Size (GB)*	*Purpose*
0	1	300	SAN OS Boot
0	2	2900	Stage Volume
1	3	1600	SQL Database LUN 1
2	4	1600	SQL Database LUN 2
3	5	1600	SQL Database LUN 3
4	6	1600	SQL Database LUN 4
5	7	1600	SQL Database LUN 5
6	8	1600	SQL Database LUN 6
7	9	1600	SQL Database LUN 7
8	10	1600	SQL Database LUN 8
9	11	1600	SQL Transaction Log LUN 1
10	12	1600	SQL Transaction Log LUN 2

To create each volume in Table 4, do the following steps, starting with Volume 1:

1. Open Hitachi Storage Navigator Modular 2.

   (1) Open a web browser and type the following in the address bar:

   `http://127.0.0.1:23015/StorageNavigatorModular/Login`

   (2) Log on using these credentials.

   User ID: `system`

   Password: `manager`

2. Open the volumes settings.

   (1) Click the **Array Name** link to open the storage system.

   (2) Expand the **Groups** heading and click the **Volumes** link.

3. Create a volume.

   (1) Click **Create Vol**.

   The **Create Volume** window displays.

   (2) For **Type**, click the **RAID Group** option.

(3) From the **RAID Group/DP Pool Number** menu, select the RAID Group number from Table 4 on page 13.

- For example, for Volume 1, click **000**.

(4) In **Vol**, type the volume number in Table 4 on page 13.

- For example, for Volume 1, type `1`.

(5) For **Capacity**, type the size in the box or select **RG ALL** for volumes that use the full capacity of the RAID group and then click the unit from the menu.

- For example, for Volume 1, type `200` in the box and then click **GB** from the menu.

(6) Leave the **Accelerate Wide Striping Mode** check box clear (unchecked).

(7) Click **OK**.

The **Create Volume** pane refreshes, populated with the new **Volume** information.

Use step 3 to create each volume in Table 4 on page 13.

Configure Host Groups

Table 5 list all default host groups used in this solution.

Table 5. Host Groups

Host HBA Name	*Host Group Description*	*Storage Port*	*Platform*
HBA1_1	SQL host group 1	0A	Windows
HBA1_2	SQL host group 2	1A	Windows
HBA1_1	SQL host group 3	0B	Windows
HBA1_2	SQL host group 4	1B	Windows

To create each host group in Table 5, do the following, starting with SQL Host Group 1.

1. Open Hitachi Storage Navigator Modular 2.

(1) Open a web browser and type the following in the address bar:

`http://127.0.0.1:23015/StorageNavigatorModular/Login`

(2) Log on using these credentials.

User ID: `system`

Password: `manager`

2. Open the host group settings.

   (1) Click the **Array Name** link to open the storage system.

   (2) Expand the **Groups** heading and click the **Host Groups** link.

3. Configure the default host group.

   (1) Click **000:G000** for port **0A**.

   (2) Click **Edit Host Group** from the top menu.

   The **Edit Host Group** screen opens.

   (3) Click **Options** tab.

   (4) Select **Windows** for **Platform** field.

   (5) Click **OK**.

Use step 3 to configure each host group listed on Table 5.

Configure Hitachi Compute Rack 220

Perform these configurations on Hitachi Compute Rack 220 before configuring the other hardware.

Configure Remote Management

Hitachi Compute Rack 220 is equipped with a baseboard management controller (BMC) mounted on the motherboard. Use this controller to access devices remotely, such as the keyboard, mouse, virtual floppy disk, and virtual CD/DVD. This remote device function allows the following:

- Operation of the system BIOS and Windows operating system
- Installation of utilities from the virtual CD/DVD

Use these steps to connect and configure the remote management for the Hitachi Compute Rack 220 server. Figure 3 shows the location of the management port used to configure the remote management. The connection can be made either by using a standard Ethernet cable and a switching hub or using a direct connection to the management port with an Ethernet cross-over cable.

Table 6 has the factory default settings for Hitachi Compute Rack 220.

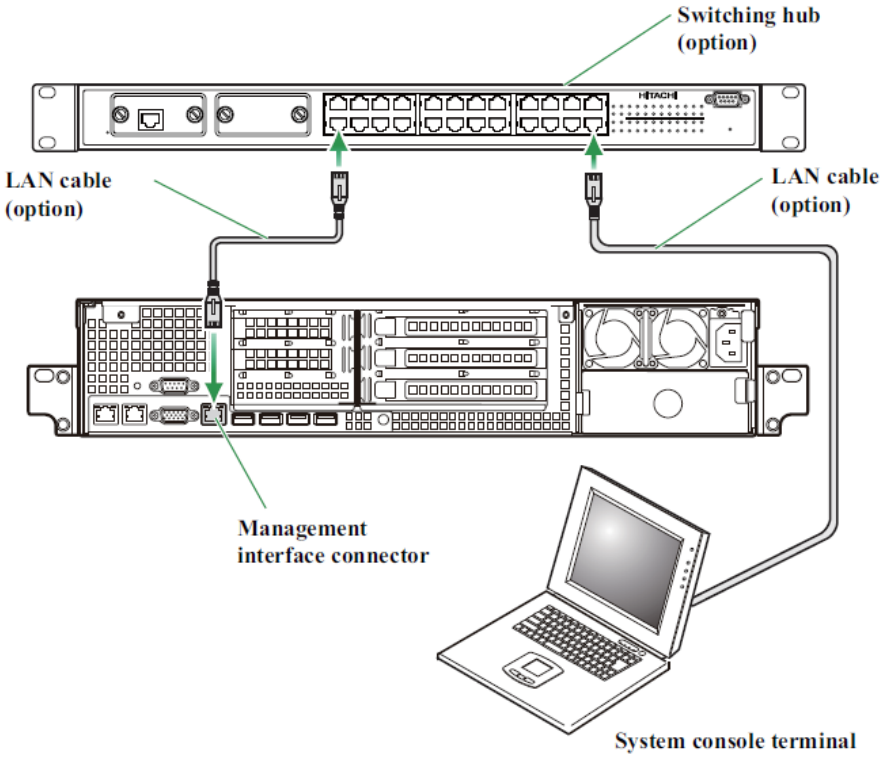

Figure 3

Table 6. Hitachi Compute Rack 220 Default Factory Settings

Item	*Factory Defaults*
IP address	192.168.100.100
Subnet mask	255.255.255.0
Default gateway	Not set
HTTP	Enabled
User Name	user01
Password	pass01

To configure the Hitachi Compute Rack 220 server, do the following:

1. Physically connect to the same switch hub.

   (1) From the **management port** on the Hitachi Compute Rack, connect a network cable to a switch hub.

   (2) From a laptop computer, connect a network cable to the switch hub.

2. On the laptop, configure a static IP of **192.168.100.101**.

3. Verify connectivity.

   - Ping the management port **192.168.100.100**.

     - If there is no reply, check your connections to the switch hub.

4. On the laptop, install the **remote console software**.

   - The software ships with the hardware. In addition, you can download it from the Hitachi Data System website.

   This software remotely manages Hitachi Compute Rack 220 and virtual devices.

5. Launch the **remote console software**.

   (1) On the desktop of the laptop, click **Remote Console**.

   (2) In **IP address**, type `192.168.100.100`.

6. Log on using the default user name and password in Table 6 on page 16.

After a successful connection, the console displays (Figure 4).

- If not, press Alt+G to display the console.

Figure 4

Access Web Console

Use the web console to do the following:

- View server and log information
- Configure networking settings
- Update BMC firmware.

To use the web console, do the following:

1. Open a web browser and type the following into the address bar:

   `http://192.168.100.100/cgi-bin/login.cgi`

2. Log on using the default user name and password (Table 6 on page 16).

Figure 5 on page 18 shows the detailed information after a successful log on.

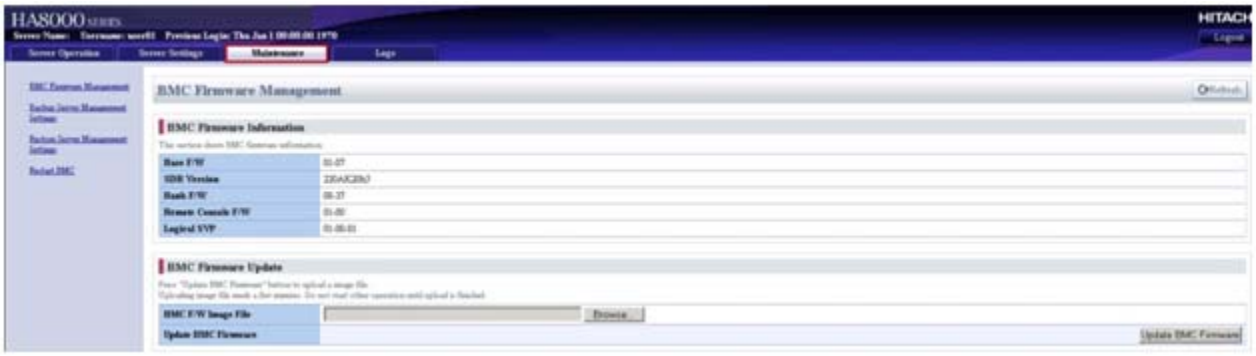

Figure 5

Configure the HBA and Server BIOS

Use this process to configure the HBA BIOS and system BIOS on Hitachi Compute Rack 220. This process is required for the SAN OS boot process to work successfully.

Do this process after racking of the server and completing all cabling.

Use the remote management console to configure the HBA and server BIOS.

To configure the HBA and server BIOS on Hitachi Compute Rack 220, do the following:

1. Install the HBA.

   (1) Verify that the power for the server is off.

   (2) Install the dual-port HBAs into the PCIe slots on the back of the server.

2. Flash the HBAs firmware.

   Flashing of the HBA firmware with the **x86 BootBIOS** boot code is required for use on the Hitachi Compute Rack 220

   - Use the **x86 BootBIOS** boot code. Get this firmware code from the **Support** section of the Emulex website.
   - For instruction on how to flash the firmware, see boot code user manual from http://www.emulex.com/downloads.

3. Configure the HBA BIOS.

   (1) Start the HBA BIOS configuration.

   i. Power the server on.

   ii. During boot up, press Ctrl+E to configure the HBA BIOS.

   iii. On the main menu, choose **Enable Boot from SAN**.

   iv. Press Esc to return to the main menu.

(2) Choose the topology.

i. On the main menu, choose **Topology** and then choose **Auto Topology: Loop First**.

ii. Press Esc to return to the main menu.

(3) Choose the link speed.

i. On the main menu, choose **Link Speed** and then choose **8 Gigaband** [*sic*].

ii. Press Esc to return to the main menu.

(4) Choose **Save** to exit the HBA BIOS screen.

4. Configure the Server BIOS.

(1) During boot up, press F2 to configure the server BIOS.

(2) On the **boot** tab, verify that the **DVD-ROM** is set to be the first device for boot priority.

(3) Press **F10** to save and to exit.

5. Make HBA 1 and Port 0 as the primary boot device.

(1) While booting the server, press Ctrl+E to enter into the HBA BIOS.

(2) On the main menu, choose **Boot Device**.

(3) From the list of saved boot devices, choose **1. Unused**.

(4) Choose a boot volume and set the boot LUN to **00** for primary boot.

(5) Press Esc to return to the main screen.

(6) Choose **Save** to exit.

Configure Storage Area Network (SAN)

Configure the storage area network (SAN) after configuring Hitachi Compute Rack 220.

HBA Direct Connection

Recommended practice is to directly connect each HBA port to use a different controller on Hitachi Unified Storage 150 for data redundancy.

Install Software

Follow these steps to install software for this solution.

Install Microsoft Windows 2008 R2

Do the following to install Microsoft Windows Server 2008 R2.

1. Place the Windows Server DVD into the DVD-ROM drive.
2. Follow the install wizard to complete the OS installation
   - Use a strong password when creating the local administrator account.
3. Install all Windows Server service packs and patches so the installation is current.
4. Depending on your Active Directory environment, join this server to an Active Directory domain before installing Microsoft SQL Server 2008 R2.

For more information, see the Microsoft TechNet article "Installing Windows Server 2008 R2."

Configure Windows Mount Point

Create and configure all LUNs from Hitachi Unified Storage 150 as windows volume or mount point, as shown on Table 7.

Table 7. Windows Volumes and Mount Points

Storage Volume	*Windows Volume*	*Mount point*	*Purpose*
LUN 0	C:		OS volume
LUN 1	D:		Staging volume
LUN 2		D:\SQL\DB_01	SQL Database Volume 01
LUN 3		D:\SQL\DB_02	SQL Database Volume 02
LUN 4		D:\SQL\DB_03	SQL Database Volume 03
LUN 5		D:\SQL\DB_04	SQL Database Volume 04
LUN 6		D:\SQL\DB_05	SQL Database Volume 05
LUN 7		D:\SQL\DB_06	SQL Database Volume 06
LUN 8		D:\SQL\DB_07	SQL Database Volume 07
LUN 9		D:\SQL\DB_08	SQL Database Volume 08
LUN 10		D:\SQL\LOG_01	SQL Transaction Log Volume 01
LUN 11		D:\SQL\LOG_02	SQL Transaction Log Volume 02

Install Hitachi Dynamic Link Manager Advanced

To install Hitachi Dynamic Link Manager Advanced, follow the instructions in the user guide that accompanies the software.

Install Microsoft SQL Server 2008 R2 Enterprise Edition

To install Microsoft SQL Server 2008 R2 Enterprise Edition, do the following.

1. Insert the SQL Server 2008 R2 DVD to launch the install wizard.
2. Click **Installation** from the menu on the left.
3. Click **New installation or add features to an existing installation**.
4. Follow the rest of the wizard to complete the installation.
5. Reboot the server so the changes take effect.
6. After the reboot, download SQL Server 2008 R2 Service Pack 1 (SP1) and update the installation.

For more information, visit the Microsoft *TechNet* Microsoft SQL Server 2008 R2

Configure Microsoft SQL Server 2008 R2

Follow these steps to configure the SQL database for a data warehouse application. It provides SQL physical file placement and additional SQL tuning parameters.

Configure SQL Server start-up parameters

The following start-up parameters are necessary for a data warehouse workload.

1. Open SQL Server Configuration Manager.
2. Click **SQL Server Services.**
3. Right-click **SQL Server (MSSQLSERVER),** and then click **Properties.**
4. On the **Advanced** tab, add this to **Startup Parameter**: **-**`E;-T1117;`
5. Click **OK** to confirm.

Restart the SQL Server service so the new settings take effect.

Additional SQL Server tuning parameters

The following tuning parameters are necessary for an optimal data warehouse workload for this solution.

1. Open Microsoft SQL Server Management Studio.

   Use **localhost** for server name.

2. From the left navigation pane, right-click the top **localhost,** and then click **Properties**. The **SQL Server Properties** dialog box opens.

3. Configure **Memory**.

   (1) Click **Memory** on the left navigation pane.

   (2) Click **Use AWE to allocate memory.**

   (3) Type the following for **Maximum server memory (in MB):** `94208`

4. Configure **Advanced**.

   (1) Click **Advanced** on the left navigation pane.

   (2) Type the following for **Max Degree for Parallelism**: `8`

5. Click **OK** to confirm.

Create a Database for a Data Warehouse Application

Table 8 shows the configuration for database, tempdb, and log files.

Table 8. Database, tempdb, and Log File Configuration Information

SQL File Type	*Database Name*	*Mount Point*	*Auto-growth Setting*	*File Type*
Database Files	Database File 1	D:\SQL\DB_01	Enable, 32 MB	Rows Data
	Database File 2	D:\SQL\DB_02	Enable, 32 MB	Rows Data
	Database File 3	D:\SQL\DB_03	Enable, 32 MB	Rows Data
	Database File 4	D:\SQL\DB_04	Enable, 32 MB	Rows Data
	Database File 5	D:\SQL\DB_05	Enable, 32 MB	Rows Data
	Database File 6	D:\SQL\DB_06	Enable, 32 MB	Rows Data
	Database File 7	D:\SQL\DB_07	Enable, 32 MB	Rows Data
	Database File 8	D:\SQL\DB_08	Enable, 32 MB	Rows Data
tempdb Files	tempdb file 1	D:\SQL\DB_01	Enable, 32 MB	Rows Data
	tempdb file 2	D:\SQL\DB_02	Enable, 32 MB	Rows Data
	tempdb file 3	D:\SQL\DB_03	Enable, 32 MB	Rows Data
	tempdb file 4	D:\SQL\DB_04	Enable, 32 MB	Rows Data
	tempdb file 5	D:\SQL\DB_05	Enable, 32 MB	Rows Data
	tempdb file 6	D:\SQL\DB_06	Enable, 32 MB	Rows Data
	tempdb file 7	D:\SQL\DB_07	Enable, 32 MB	Rows Data
	tempdb file 8	D:\SQL\DB_08	Enable, 32 MB	Rows Data
Transaction Log	Database Transaction Log	D:\SQL\LOG_01	Enable, 32 MB	Log
	Tempdb transaction Log	D:\SQL\LOG_02	Enable, 32 MB	Log

Note—The basis for making the recommended auto-growth settings is the tested database for this solution. Actual settings for your deployment may vary based on the database capacity and business requirements for each data warehousing solution.

To create a database for a data warehouse application, do the following.

1. Open Microsoft SQL Server Management Studio.

   Use **localhost** for server name.

2. From the left navigation pane, right-click the Database folder and then click **New Database**. The new database screen opens.

3. On the **Database name** field, provide a name for the new database that is being created.

4. On the **Database files**, create, configure, and place eight database files into each of the DB volumes listed in Table 4 on page 13.

   - Configure each database file **Initial Size (MB)** and **Autogrowth** settings, as shown in Table 8.

5. Configure the placement of **Log files** to the Log Volumes. See Table 7 on page 20.

6. Click **OK** to confirm.

Configure Database, tempdb and Log Files

To configure the SQL Server tempdb setting and file placement, do the following.

1. Open Microsoft SQL Server Management Studio.

   Use **localhost** for server name.

2. From the left navigation pane, click **Database** and then click **System Database**.

3. Right-click the **tempdb database** and click **Property**. The dialog box for the tempdb database properties opens.

4. Click **Files** from the left navigation.

5. In the **Database files**, create and configure eight database files, placing each file into its DB volume. See Table 8 on page 22.

6. Set **Initial Size (MB)** and **Autogrowth** for each database file. See Table 8.

7. Configure the **Log files** placement in the Log Volumes, using Table 8.

8. Click **OK** to confirm.

For More Information

Hitachi Data Systems Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the Hitachi Data Systems Global Services website.

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the Hitachi Data Systems Corporate Resources website. Click the **Product Demos** tab for a list of available recorded demonstrations.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the Hitachi Data Systems Services Education website.

For more information about Hitachi products and services, contact your sales representative or channel partner or visit the Hitachi Data Systems website.

Hitachi Data Systems Corporation

Corporate Headquarters
2845 Lafayette Street, Santa Clara, California 95050-2639 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