

Virtualizing Microsoft[®] SQL Server 2008 Using VMware vSphere 4 on the Hitachi Adaptable Modular Storage 2000 Family

Deployment Guide

By Eduardo Freitas and Bob Ingram

November 2009



Summary

Deploying a virtualized Microsoft® SQL Server 2008 environment with VMware vSphere 4 can be a complex endeavor without proper planning and deployment guidance. However, when effectively implemented, the benefits are significant, including reduced storage and server footprints, reduced power and cooling costs, simplified management of the SQL Server infrastructure, high availability and easily scalable storage.

This white paper describes the planning and deployment process using a building block architecture to help large and enterprise deployments achieve these critical business objectives. It provides planning, installation, and configuration procedures and related information required to successfully deploy SQL Server 2008 instances in a virtualized environment using VMware vSphere 4 configurations on the Hitachi Adaptable Modular Storage 2000 family.

For best results use Acrobat Reader 8.0.



Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an e-mail message to SolutionLab@hds.com. Be sure to include the title of this white paper in your e-mail message

Table of Contents

Solution Overview	1
Hardware Resources	3
Hitachi Adaptable Modular Storage 2000 Family	3
Servers	3
Storage Area Network	4
Software Resources	5
Planning for Deployment	6
Deployment Scenarios	6
Sizing ESX Hosts	6
Sizing SQL Server 2008 Virtual Machines	7
Planning Storage Performance and Capacity	8
Deploying the Solution	9
Configuring the SAN	9
Storage Configuration	10
Using Hitachi Dynamic Provisioning Software	18
Virtual Server Installation	22
Deploying SQL Server Virtual Machines	23
Validating Your Deployment	25
Monitoring and Managing the Solution	25
Managing a Growing Environment	25
VMware Product Features	26
Hitachi Performance Monitor Feature	26
Hitachi Tuning Manager Software	27

Virtualizing SQL Server 2008 Using vSphere 4 on the Hitachi Adaptable Modular Storage 2000 Family

Deployment Guide

Deploying a virtualized Microsoft® SQL Server 2008 environment with VMware vSphere 4 can be a complex endeavor without proper planning and deployment guidance. However, when effectively implemented, the benefits of virtualizing a SQL Server 2008 instances using VMware vSphere 4 and the Hitachi Adaptable Modular Storage 2000 family outweigh the complexity in many ways, and include the following:

- Reduced storage and server footprints and power and cooling costs through consolidation
- Simplified management of SQL Server infrastructure
- High availability and easily scalable storage deployments with predictable performance and capacity

This white paper describes the planning and deployment process using a building block architecture to help large and enterprise deployments achieve these critical business objectives. It provides planning, installation, and configuration procedures and related information required to successfully deploy SQL Server 2008 instances in a virtualized environment using VMware vSphere 4 configurations on the Hitachi Adaptable Modular Storage 2000 family. It describes the prerequisites, components and procedures needed to successfully complete the following tasks:

- Locate all required solution components and related documentation
- Install and configure the required solution components
- Validate and manage the deployed solution

The white paper focuses on planning and deploying a SQL Server 2008 instance in a virtual machine and is intended for use by IT and database administrators responsible for SQL Server instances, server virtualization or storage. It assumes familiarity with VMware vSphere 4 and vCenter, Hitachi Storage Navigator Modular 2 software, Microsoft Windows 2008 and Microsoft SQL Server 2008.

Although the documented solution in this paper used the Hitachi Adaptable Modular Storage 2300 for testing, any member of the Hitachi Adaptable Modular Storage 2000 family provides a reliable, flexible and cost-effective storage platform for supporting demanding applications like SQL Server 2008.

Solution Overview

The building block architectures described in this white paper were developed with ease of management and scalability in mind. The architectures consist of a vSphere virtual machine (VM) running Windows Server 2008 and SQL Server 2008 with Service Pack 1 with storage hosted on the Hitachi Adaptable Modular Storage 2300. VMware recommends deploying virtual machines on shared storage like the 2300 to take advantage of VMware features such as VMotion, VMware High Availability (HA), and VMware Distributed Resource Scheduler (DRS).

Logical units (LUs) from the 2300 are allocated to the ESX hosts and formatted as VMFS volumes in which virtual disks can be created. The virtual disks (vDisks) are presented to the Windows 2008 operating system and can be partitioned and used in file systems for the SQL Server databases and transaction logs.


The solution building block described in this white paper supports a 500GB database as a starting point. This configuration meets or exceeds the Microsoft guidelines for acceptable response time for OLTP databases,

which is at most 20ms for databases reads, database writes and log reads, and at most 5ms for log writes. Each SQL Server instance was appropriately configured to provide enough CPU and memory resources for the capacity it was configured with and a load that maintains at least 50 percent CPU busy rate while maintaining the acceptable response times for each of its LUs.

The following figures show the underlying storage configurations for the 500GB database building block for the traditional storage configuration, not using the dynamic provisioning approach.

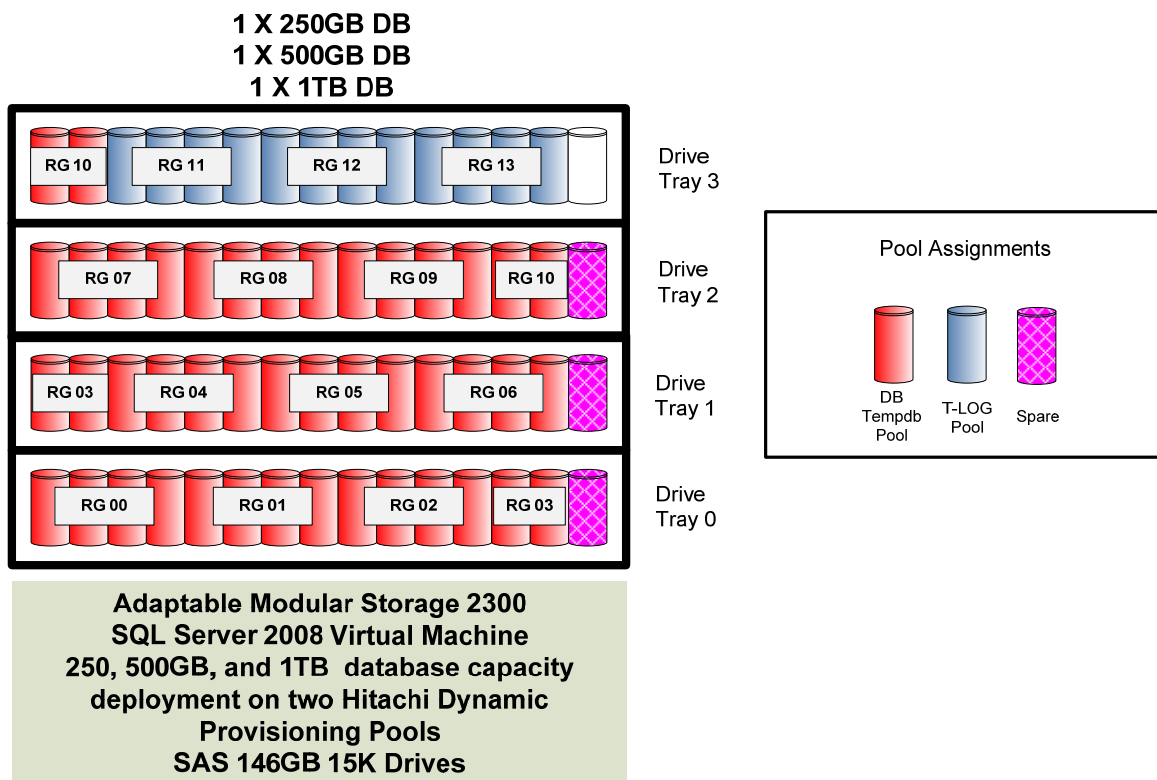
Figure 1 shows the building block for the 500GB capacity environment.

Figure 1. Standard Provisioned 500GB Storage Building Block



For nontraditional storage configurations utilizing Hitachi Dynamic Provisioning software, create two Dynamic Provisioning pools, one for the SQL Server database, and tempdb volumes and a separate one for the transaction logs volumes. This gives you the best performance and minimizes hot spots across the RAID groups while maintaining a protection level from the separation between the location of the database and transaction log files. Nevertheless, dedicated volumes within the pool simplify backup operations due to their granularity and also provide an easier way to monitor the capacity utilization of each of the volume types. For example, if you are creating a Dynamic Provisioning pool for an environment containing three SQL Server instances having a 250GB, 500GB, and 1TB capacity, use a total of 14 RAID-1+0 2+2 groups, 11 of which are used for the database and tempdb volumes, and 3 of which are used for the transaction log volumes. Figure 2 shows the example architecture that is built in the same manner as the building block architecture. Note that all RAID-1+0 2+2 groups shown in Figure 2 are assigned to the Dynamic Provisioning pools.

Figure 2. Dynamic Provisioning Environment of a 250GB, 500GB and 1TB Database Deployment



For more information about building block architecture and Hitachi Dynamic Provisioning software, see the [Virtualizing Microsoft® SQL Server Using VMware vSphere 4 on the Hitachi Adaptable Modular Storage 2000 Family Best Practices Guide](#) white paper.

Hardware Resources

The following sections describe the hardware components used to deploy the solution described in this white paper.

Hitachi Adaptable Modular Storage 2000 Family

This solution was tested on a Hitachi Modular Storage 2300. However, any of the Hitachi Adaptable Modular Storage 2000 family storage systems can be utilized. Determining which member of the 2000 family is best for you requires evaluating your overall capacity and connectivity requirements.

Hitachi Data Systems testing used a Hitachi Adaptable Modular Storage 2300 with 146GB 15K RP SAS disks, 8GB cache per controller and four ports, two per controller.

Servers

Use two enterprise class level servers when deploying this SQL Server solution in high availability environments. The servers must have enough processor and memory resources to operate the SQL Server environments and meet vSphere operating requirements. They both must have the same hardware architecture and configuration and processor and memory resources to accommodate a failover situation. Best practice is to not over-commit the ESX host resources (CPU or memory), as it can lead to a bottle neck of resources during peak database utilization times.

For more information about vSphere server hardware requirements, see the [VMware Hardware Compatibility List](#).

Storage Area Network

For this solution, Hitachi Data Systems recommends connecting the vSphere servers and the Adaptable Modular Storage 2300 to two Fibre Channel switches for redundancy and high availability purposes. Another option is to deploy the environment on an enterprise level switch that contains different blades for connections.

In addition, Hitachi Data Systems recommends configuring a minimum of two redundant paths from the ESX host to the 2300. Each ESX host requires at least two host bus adapters (HBAs) for high availability. VMware's round-robin load balancing algorithm automatically selects a path by rotating through all available paths, thus balancing the load across all available paths, optimizing IOPS and response time.

Figure 3 illustrates the storage area network configuration.

Figure 3. Minimum Storage Area Network Configuration




Table 1 lists the hardware used in testing this architecture.

Table 1. Tested Deployment Hardware Resources

<i>Hardware</i>	<i>Configuration</i>
Hitachi Adaptable Modular Storage 2300	Microcode 0870H* 2 controllers 4 disk trays 15 146GB 15K RPM SAS disks per tray 16GB cache per controller
Two Brocade 5100 SAN switches	FOS 6.21* 16 4Gb Fibre Channel ports
Two Dell R905 servers	BIOS firmware 4.0.3** 4 quad core AMD Opteron 1.9GHz processors 128GB memory 2 Emulex LPe11002 4Gb Fibre Channel HBAs

*Use the latest or the most recent Hitachi supported levels of Microcode and FOS levels

**Use the latest available or VMware supported BIOS firmware level.

Software Resources

Table 2 describes the software used in the tested deployment described in this white paper.

Table 2. Tested Deployment Software Resources

<i>Software</i>	<i>Version</i>
ESX	4.0.0, build 140815
vCenter	4.0, build 140742
Windows Server	2008 Enterprise edition, Service Pack 2
SQL Server 2008	2008 Enterprise edition, Service Pack 1
Hitachi Storage Navigator Modular 2	7.0*
Hitachi Device Manager	6.2.0-00
Hitachi Tuning Manager	6.2.0-00

*Note: To use Hitachi Dynamic Provisioning software in this solution, the license for the Hitachi Dynamic Provisioning software is also required. The license can be installed using Hitachi Storage Navigator Modular 2 software's GUI or command-line interface.

For more information about the software used to deploy this solution, see the Virtualizing SQL Server on the Hitachi Adaptable Modular Storage 2000 Family Using ESX Server Best Practices Guide, or visit the Hitachi Data Systems [Storage Software](#) Web site.

Planning for Deployment

Like any other SQL Server 2008 environment, the deployment of virtualized SQL Server instances on vSphere requires careful planning and proper testing prior to becoming part of your production environment. The following sections describe the planning process for deploying a virtualized instance of the SQL Server 2008 on vSphere 4.

Backup schemes and log shipping are vital to an organization's SQL Server environment, however these tasks fall outside of the scope of this paper. For more information, see the [Hitachi Adaptable Modular Storage 2000 Family Best Practices with Microsoft® SQL Server For OLTP Applications](#) white paper or the [Backing Up and Restoring Databases in SQL Server](#) Books Online Web site.

Deployment Scenarios

This deployment guide concentrates on deploying SQL Server 2008 on vSphere for on a 500GB capacity building block. While this building block approach supports different deployments, the information in this white paper enables you to deploy any of them by repeating the steps described in this guide

Sizing ESX Hosts

Before you can establish the specifications for your ESX hosts, you must first determine how many virtual machines will run on each ESX host. You must also determine the virtual CPU, memory and capacity requirements for virtual machines. You also need to consider whether the ESX host will be part of a cluster and whether features like VMotion will be used.

For more information, see the Virtualizing SQL Server on the Hitachi Adaptable Modular Storage 2000 Family Using ESX Server Best Practices Guide.

Virtual CPUs

A virtual environment contains both virtual CPUs (vCPUs) and physical CPUs. Virtual CPUs are the processors within the virtual machines, and they equate to one physical processor core. The virtual machine monitor is responsible for virtualizing the CPUs. VMware recommends ensuring that the total number of vCPUs assigned to all the virtual machines is equal to or less than the total number of cores on the physical ESX host. Keep in mind that the ESX service console requires one CPU and at least 400MB of memory.

vSphere Memory Configuration Guidelines


vSphere manages memory at the hypervisor level, enabling such features as memory sharing across VMs that have the same guest operating systems, memory over-commitment and a memory balloon technique that allows virtual machines that are not using all of their allocated memory to give memory to virtual machines that require more.

When sizing a ESX host for memory, consider two additional factors beyond what is needed for SQL Server:

- **ESX host memory overhead** — Overhead for the service console, the VMkernel and each virtual machine
- **Virtual machine memory settings** — Overhead for the ESX host, also called guest memory, which can be broken down into the following parts:
 - **Configured memory** — Memory size assigned to the virtual machine
 - **Touched memory** — Amount of virtually allocated memory the virtual machine is using
 - **Reserved memory** — Portion of configured memory vSphere reserves for the virtual machine, taking into account overhead
 - **Swappable memory** — Portion of configured memory beyond reserved, which can be reclaimed by the balloon driver or by vSphere swapping

Figure 4 illustrates virtual machine memory settings.

Figure 4. Virtual Machine Memory Settings



Sizing SQL Server 2008 Virtual Machines

The two main specifications that need to be determined when sizing SQL Server 2008 servers are the number of processor cores and memory. While the ESX host has hardware requirements of its own, you must ensure that they CPU and memory requirements for the SQL Server virtual machines match the recommended configurations as per Microsoft requirements. For a list of the hardware and software requirements for installing SQL Server 2008, see the [Hardware and Software Requirements for Installing SQL Server 2008](#) books online article. Also keep in mind that different SQL Server editions support a maximum number of CPUs. For more information, see the [Maximum Number of Processors Supported by the Editions of SQL Server](#) books online article.

SQL Server Processor and Memory Considerations

Establishing the number of processors and amount of memory a given SQL Server needs depends on many factors such as SQL Server 2008 edition, backup strategy, the utilization of data compression, reporting services, and most importantly user workload for the instance databases. While Hitachi Data Systems tests show that the building block architecture described in this white paper fully utilizes the virtual machine resources, it is important to first deploy your environment as a test platform. One of the key benefits of virtualizing SQL Server instances is the capability and ease of adding resources such as CPUs and memory to a virtual machine when needed.

From a SQL Server instance configuration perspective, be sure to establish a number of tempdb files that matches the number of vCPUs assigned to the virtual machine as per Microsoft's best practice recommendations. Keep in mind that unlike a CPU, which contains many cores, a single vCPU only provides a single ESX host CPU core to a virtual machine.

For example, for this solution's deployed architecture, two vCPUs are allocated to the SQL Server virtual machine, that means a total of two tempdb files need to be made available to this SQL Server instance. Through SQL Server Management studio, a modification of the existing tempd file location and the addition of one more tempd file can be made by executing the following SQL statements.

```
use tempdb
go

-- 2 vCPU Machines --

alter database tempdb
modify file (name='tempdev', filename='T:\SQL\DATA\datatempdb1.mdf', size=8192 KB, FILEGROWTH=10MB)

alter database tempdb
add file (name='tempdev2', filename='T:\SQL\DATA\datatempdb2.ndf', size=8192 KB, FILEGROWTH=10MB)

sp_helpfile
go
```

The last statement provides an output of the current tempdb files configured on the SQL Server instance.

Note that the files are now located on a mount point that was created from a LU that is assigned to the ESX host to be used as the datastore for a given SQL virtual machine's tempdb.

For more information about tempdb and how to optimize it, see the [Optimizing tempdb Performance](#) Books Online article.

Planning Storage Performance and Capacity

Deploying SQL Server 2008 using a building block approach allows you to easily manage and scale your environment. Additional VMs are easy to deploy using templates, and storage can be provisioned on a per RAID group basis. This white paper provides step-by-step instructions explaining how to implement the 500GB building block architecture that can be scaled up according to your performance and capacity needs. Keep in mind that from a performance perspective, scaling up might require a redistribution of your database files and additional hardware resources for your virtual SQL Server implementation.

The architecture described in this guide utilizes a RAID-1+0 6+6 group that accounts for a 20 percent unexpected growth space as per Microsoft's best practice and space to allocate a LU to be used by tempdb. The transaction logs utilize a RAID-1+0 2+2 group that also accounts for a 20 percent growth space. For more information, see the Virtualizing SQL Server on the Hitachi Adaptable Modular Storage 2000 Family Using ESX Server Best Practices Guide.

The storage used in this solution is primarily for SQL Server 2008 databases, tempdb and transaction logs, although some storage is also needed for guest operating systems. Sizing storage for a virtual server is no different than for a physical server. Follow Hitachi Data Systems and Microsoft best practices for planning storage capacity for SQL Server databases and logs. Compute the following:

- Overall capacity needed for databases and transaction logs including planned growth for the SQL Server instances
- IOPS needed for SQL Server databases and transaction logs
- SQL Server instance host CPU and memory requirements for the expected database workloads

Be sure to account for overhead when sizing the databases and logs. A typical overhead measure is to keep provision approximately 20 percent more storage than what is expected to be used for a given database implementation. After the underlying storage is sized for IOPS and capacity, allocate the LUs to the ESX hosts. Keep the following limitations in mind:

- You can allocate a maximum of 256 LUs per ESX host.
- You can create a maximum of 60 vDisks per VM.

Deploying the Solution

This section describes how to configure the storage, SAN and virtual machines to build a successful and reliable virtual SQL Server environment using Hitachi Adaptable Modular Storage 2300. While the building block being used on this deployment is for a 500GB capacity, additional blocks of the same architecture can be added to the environment to either create the capacity of a given SQL instance or to create a new instance with the same capacity requirements.

Table 3 lists the building block resources needed to scale the architecture in 500GB increments.

Table 3. 500GB Building Block Resources Using Standard Provisioned RAID Groups

<i>Resource</i>	<i>Details</i>
Virtual machines	1 32GB memory 2 vCPU
Storage systems	1 Hitachi Adaptable Modular Storage 2300 23 x 146GB 15K RPM SAS disks Note: 2 disks are used as hot spares
RAID groups	1 x RAID-1+0 (6D+6D) 1 x RAID-1+0 (2D+2D) 1 x RAID-1+0 (4D+1P)
LUs	1 x 625GB on RAID-1+0 (6D+6D) for database files 1 x 62.5GB on RAID-1+0 (6D+6D) for tempdb files 1 x 160GB on RAID-1+0 (2D+2D) for transaction logs 1 x 150GB on RAID-5 (4D+1P) for VM OS*

* The remainder of the space within the RAID-5 (4D+1P) group was used for additional virtual machines that were hosted on the same ESX host.

Configuring the SAN

This solution requires two or more ESX hosts to form a highly available cluster. Configuring the SAN in an ESX environment requires careful consideration of all ESX hosts in the cluster. Configure the SAN zones so that the ports on the ESX hosts running the SQL Server instances has visibility to four ports of the 2000 family storage system: two ports from each storage system's controller. While only two ports from the ESX hosts are used in this environment, more ports can be configured to provide higher availability. See Figure 3 for an illustration of the connection between the ESX hosts and the storage system.

Note that while four dedicated ports from the host and four from the storage system were utilized for each host to storage connection on the test, two ports from the host and two from the storage system can also be used without reaching the IOPS capability of any port and therefore not creating any type of bottleneck in the environment.

Ensure that both the storage and host Fibre Channel ports are configured to point-to-point connection. This is the connection type that must be used when a Fibre Channel switch is used in your environment

Table 4 lists the connections between the primary ESX host and the storage system ports.

Table 4. Path Configuration for a Single ESX Host

<i>ESX Host</i>	<i>Host HBA Number / ESX Port Name</i>	<i>Switch</i>	<i>Zone Name</i>	<i>Storage System Port</i>	<i>Storage System Host Group</i>
ESX1	HBA 1 port 1 vmHBA 1	Brocade 5100 - 1	ESX_1_HBA1_1_AMS2300_0A	0A	ESX_1_HBA1_1
ESX1	HBA 1 port 2 vmHBA 2	Brocade 5100 - 2	ESX_1_HBA1_2_AMS2300_1A	1A	ESX_1_HBA1_2
ESX1	HBA 2 port 1 vmHBA 3	Brocade 5100 - 1	ESX_1_HBA2_1_AMS2300_0B	0B	ESX_1_HBA2_1
ESX1	HBA 2 port 2 vmHBA 4	Brocade 5100 - 2	ESX_1_HBA2_2_AMS2300_1B	1B	ESX_1_HBA2_2

Storage Configuration

The Hitachi Adaptable Modular Storage 2000 family allows for easy configuration and deployment of storage solutions through the use of the Hitachi Storage Navigator Modular 2 software and Hitachi Device Manager. This white paper provides step-by-step instructions explaining how to provision and manage the storage for your SQL Server environment using the Storage Navigator Modular 2 software GUI. A command-line interface for Storage Navigator Modular 2 software is also available.

This solution uses RAID-1+0 for both database and transaction logs. When deploying the environment using standard LUs, keep these two types of storage in separate RAID groups to prevent contention and latency problems.

This solution focuses on provisioning the SQL Server portion of the storage, however, the same procedure can be used when deploying the virtual machines' operating system storage.

Creating RAID Groups

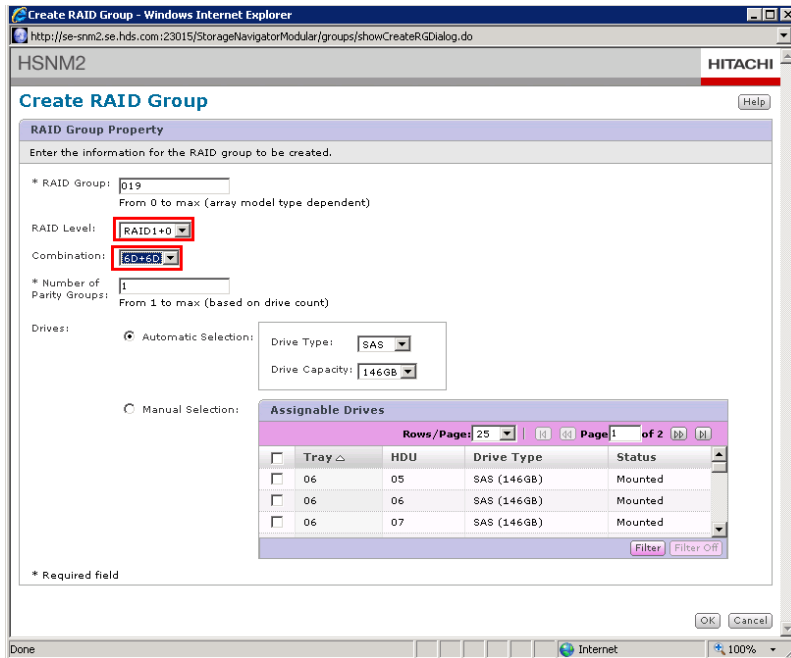
Follow these steps to create RAID groups using Hitachi Storage Navigator Modular 2 software:

1. Log in to the Storage Navigator Modular 2 instance on your data center.
Ensure that you have modify privileges on Storage Navigator Modular 2 software.
2. Click the **Array Name** link to open the storage system to be used for the SQL Server environment.
3. Expand the **Groups** heading in the storage system's pane window and click on **Logical Units** link.

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

4. Click the **RAID Groups** tab and click the **Create RG** button.
5. Select **RAID1+0** from the **RAID Level** drop-down menu and **6D+6D** from the **Combination** drop-down menu.

This is the RAID group for database and tempdb LUs.



6. Leave the **number of parity groups** at 1.

This is a Hitachi Data Systems recommendation. Increasing the number of parity groups above the default of 1 allows the creation of large LUs across multiple RAID groups. Hitachi Data Systems recommends against the creation of LUs that span across multiple RAID groups. Instead, use Hitachi Dynamic Provisioning software to create LUs that automatically span across multiple hard drives.

7. Select the **Automatic Selection** radio button.

If different types of drives are installed in the storage system (for example SATA and SAS, or different capacities), be sure to select the correct type so that the Storage Navigator Modular 2 software can select the right type of hard drive.

Enabling automatic selection means that the Storage Navigator Modular 2 software selects the next available drives of the type shown in the box. Hitachi Data Systems recommends enabling this setting.

8. Click **OK**.

A message displays stating that the RAID group was successfully created.

Repeat these steps to create the RAID group that will host the transaction log LUs as described in Table 3.


Creating LUs

Follow these steps to create LUs using Hitachi Storage Navigator Modular 2 software.

1. Expand the **Groups** heading in the storage systems' pane window and click the **Logical Units** link.
2. Select the **RAID Groups** tab and click the RAID group where you want to create the LU.

The Summary window displays in the right pane.

3. Click the **Create LU** button.



4. Change the **Capacity** information from RG ALL to GB, and type the desired LU size.

In this example, the first LU to be created on the RAID group needs 625GB. Remember that while this is a 500GB database, 20 percent more capacity is being allocated for future and unplanned growth as per Microsoft's best practice.

5. Click **OK**.

A message displays stating that the logical unit was created successfully.

The newly created LUs begin a quick formatting process that makes them available immediately for use.

Repeat these steps to create a tempdb LU within the same RAID group and then to create a transaction log LU in its own RAID group..

Configuring Fibre Channel Port Settings


In preparation for this step, all ESX HBAs need to be connected to a Fibre Channel switch and appropriately zoned to the storage system ports. Contact your SAN administrator to ensure that the proper zoning configuration is in place in your environment. It is important to ensure the correct port is being edited, as changes to the port's topology setting cause I/O interruption between any host that might be connected to the port.

To ensure that your storage Fibre Channel ports are configured for the correct **Topology** configuration, follow these steps:

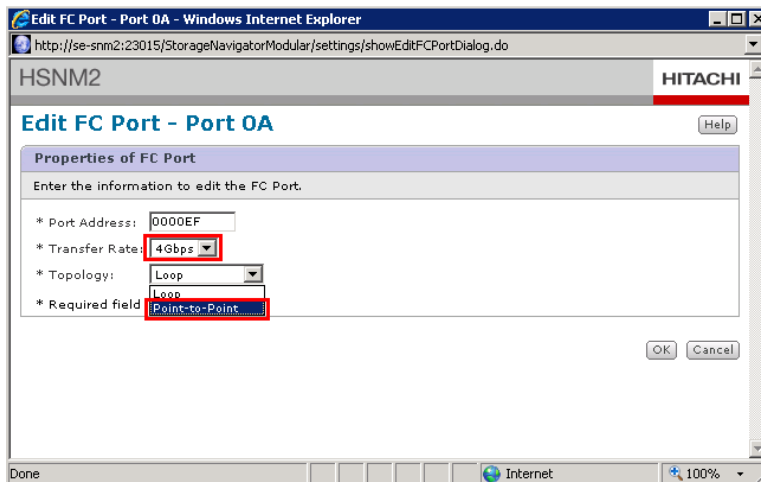
1. Expand the **Settings** heading and click on the **FC Settings** link.

Notice the **Topology** column for the ports that were zoned to connect to the ESX hosts on the SAN. If any of them are not configured to **Point-to-Point**, you must change their configuration.

2. Click the **Edit FC Port** button.



3. Choose **Point-to-Point** from the Topology drop-down menu and click **OK**.



A disclaimer displays indicating that the change will cause I/O interruption between any host that is connected to the port at the moment.

4. Click the **Confirm** button and wait a couple seconds for the change to take place.

A message displays stating it is necessary to refresh Storage Navigator Modular 2 software information to see the change. Note that it might also be necessary to reset the ESX host's ports before they can connect to the storage system.

After the connection between the storage system and the ESX host is established, the FC Settings screen shows all modified ports with a **LinkUp(F_Port Connectec)** status.

Enabling Host Group Security

Enabling Host Group Security enabled you to associate specific WWNs to a given port as well as a set of LUs within that port. This enables you to have more granularity in cases where more than a single host is connected to the storage port.

Follow these steps to enable Host Group Security using Hitachi Storage Navigator Modular 2 software:

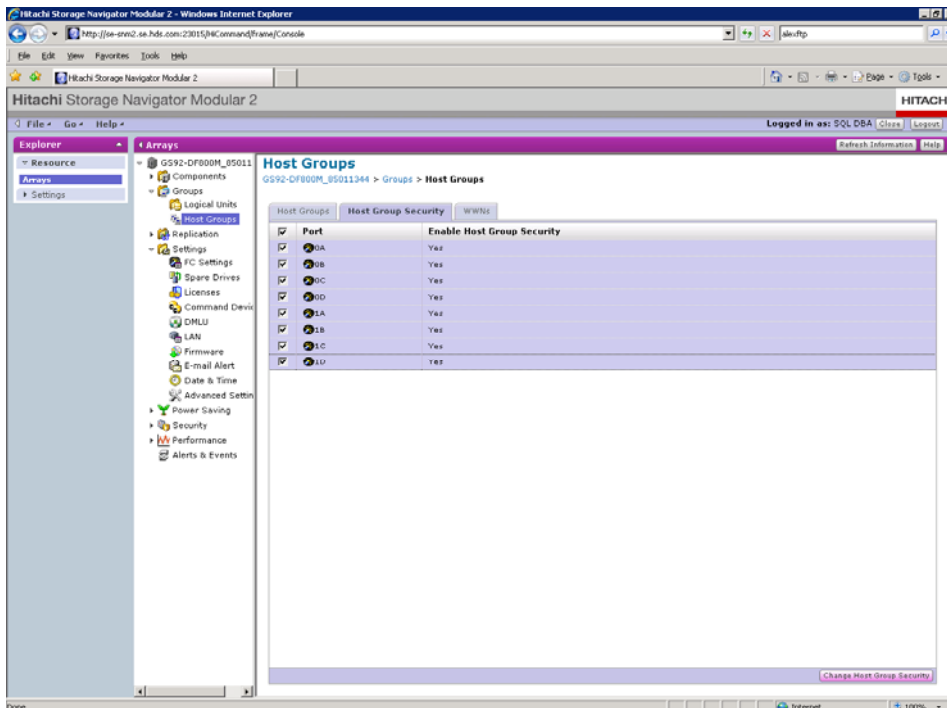
1. Expand the **Groups** heading and click the **Host Groups** link.
2. Click the **Host Group Security** tab and ensure that port security is enabled on all ports.

This ensures that zoning to specific ports on the host can be created on each of the storage ports.

If the host security is not enabled, select all ports to be used in the environment and click the **Change Host Group Security** button.

3. Select the **Yes** checkbox and click **OK**.

The **Host Security Group** setting is now enabled on all selected ports.



Assigning Nicknames to HBA's WWNs

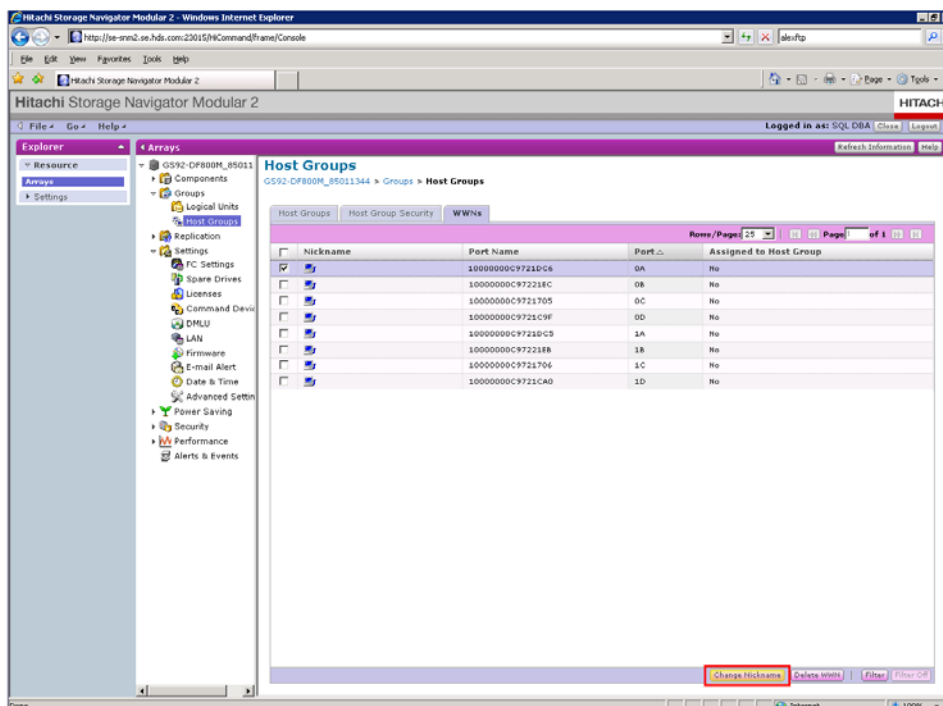
This procedure is optional, however, assigning a nickname to ESX host HBA's WWNs can facilitate on identifying each host and port to ease documenting and troubleshooting your environment. Note that this step assumes that the necessary connections and zones exist on the SAN and that the correct settings are configured on storage ports and host HBAs.

Follow these steps to assign a nickname to a HBA's WWN using Hitachi Storage Navigator Modular 2 software:

1. Expand the **Groups** heading and click the **Host Groups** link.
2. Click the **WWNs** tab.

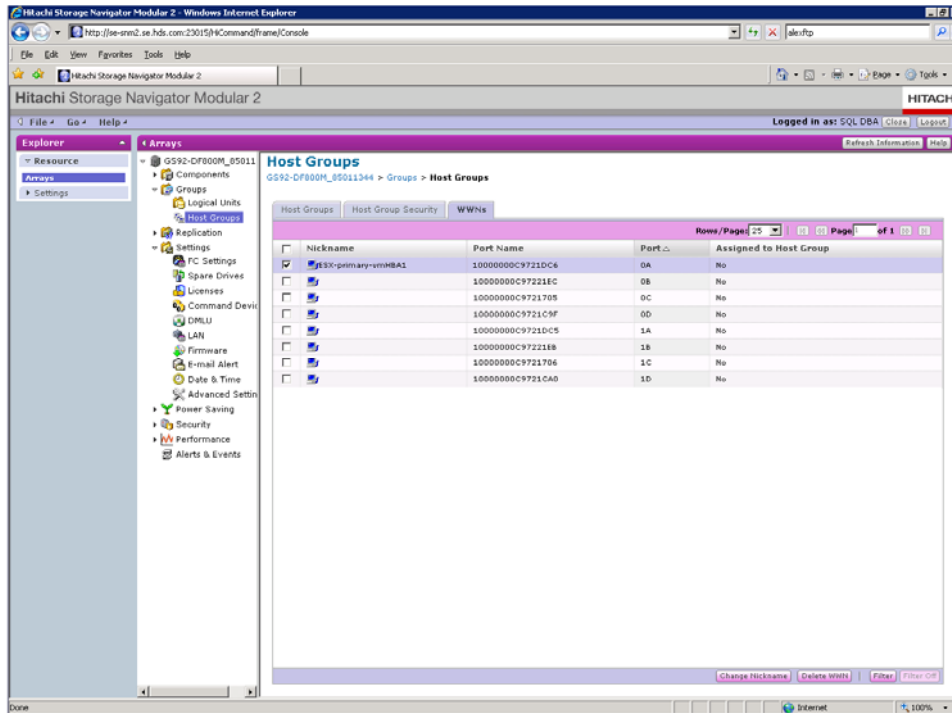
This tab shows all previously zoned WWNs with each WWN assigned to a Fibre Channel port.

3. Select a port by checking its box in the far left column.
4. Click the **Change Nickname** button.



5. Enter the desired nickname of the WWN being edited and click **OK**.
A message displays stating the edition of the WWN nickname is complete.
6. Click the **Close** button.

- The Storage Navigator Modular 2 software screen automatically refreshes showing the new nickname for the WWN.




Repeat this step for the ESX host's remaining WWNs.

Assigning LUs to a Host Group

Follow these steps to assign LUs to a host group using Hitachi Storage Navigator Modular 2 software:

1. Expand the **Groups** heading and click the **Host Groups** link.
2. Click the **Create Host Group** button.



The **Host Group No.** field can remain the same because Storage Navigator Modular 2 software automatically selects the next available number.

3. Change the **Name** of the host group to something that identifies the hosts connected to it.

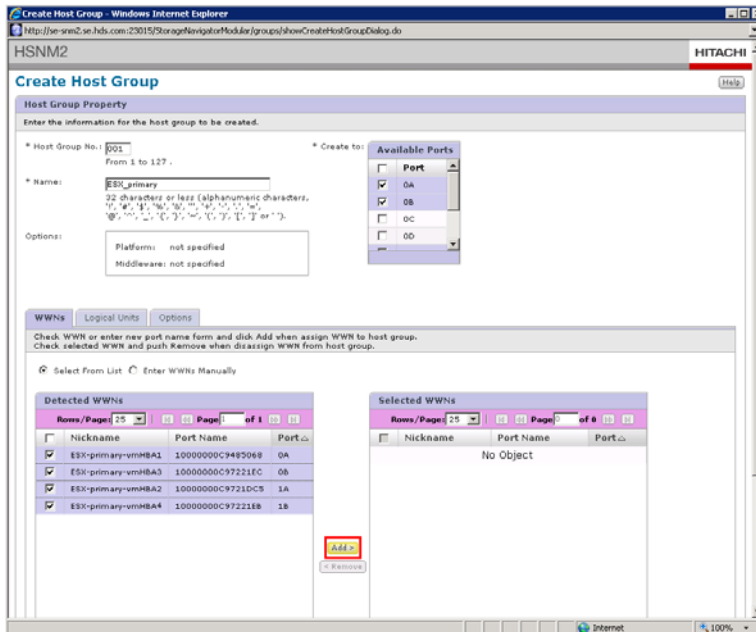
For example, you might use ESX primary for your primary ESX host.

4. Select all of the Fibre Channel ports that are connected to the ESX host.

This ensures that multiple paths to the LUs to be used are available both for high availability and failover purposes.

Note that if the HBA ports are not detected by the storage system, their WWNs are not shown in the **Detected WWNs** window. In that case, it might be necessary to troubleshoot the environment to ensure that the proper settings are set on both the storage ports and host HBAs.

- Select the World Wide Names for the host group and click **Add**.



The screen automatically refreshes and shows the added WWNs in the **Selected WWNs** column.

- Click the **Logical Units** tab.
- Select the LUs to be added to the host group and click **Add**.

The screen automatically refreshes and shows the added LUs in the **Assigned Logical Units** column.

- Click the **Options** tab.
- Select **VMware** from the **Platform** drop-down menu and click **OK**.

Repeat this procedure for other ESX hosts being added to the storage cluster by selecting the appropriate ports, host WWNs and LUs.

Using Hitachi Dynamic Provisioning Software

The 2000 family can be configured to use Hitachi Dynamic Provisioning LUs instead of standard provisioned LUs. In this configuration, all data is distributed across all of the hard disk drives (HDDs) in the Dynamic Provisioning pool. This helps prevent contention for I/O on heavily used LUs by distributing the usage across all HDDs.

Dynamic provisioning requires the following components:

- Hitachi Adaptable Modular Storage microcode version 0870 or higher
- Storage Navigator Modular 2 software version 7.0 or higher
- Hitachi Dynamic Provisioning software license key

Hitachi Dynamic Provisioning software enables you to create a Dynamic Provisioning pool that is made up of one or more RAID groups. A Dynamic Provisioning LU does not consume space from the Dynamic Provisioning pool until the host writes to the Dynamic Provisioning LU.

Hitachi Storage Navigator Modular 2 software default alert settings for Hitachi Dynamic Provisioning match Hitachi Data Systems recommended minimal settings for a given Dynamic Provisioning pool, as follows:

- **DP Pool Consumed Capacity Alert** — This alert monitors the use rate of a Dynamic Provisioning pool and warns when the capacity is nearly exhausted; issues one of two alerts:
- **Early Alert** — This is an initial warning. The default value is 40 percent. Hitachi Data Systems recommends that you consider the time it might require to order and set up additional drives.
- **Depletion Alert** — This warning signals that available capacity is nearly exhausted. The storage system provides warnings every six hours if this state continues. The range is from 1 to 99 percent. If the consumed rate exceeded these alert levels, the storage system can report by either the Simple Network Management Protocol (SNMP) Agent or an e-mail alert feature.
- **Over Provisioning Threshold** — This warning is provided when the total capacity of LUs in a Dynamic Provisioning pool is larger than a Dynamic Provisioning pool capacity according to the following ratio:

$$\text{(Total Virtual Capacity)} \div \text{(Total Pool Capacity)}$$

These limits help reduce the risk of excessive over-provisioning. If the ratio is greater than 1.0 (or greater than 100 percent), the system is over-provisioned and pool capacity is insufficient to support all of the defined virtual volumes' capacity.

For example, a ratio of 1.4 (140 percent) defines a "comfort limit" ratio by defining more virtual capacity than the current pool capacity can currently support.

The following sections provide instructions on how to create each of the Hitachi Dynamic Provisioning components.


Creating a Dynamic Provisioning Pool

Follow these steps to create a Dynamic Provisioning pool:

1. Log in to the Storage Navigator Modular 2 software instance on your data center.
You must have modify privileges for Storage Navigator Modular 2 software.
2. Click the **Array Name** link to open the storage system to be used for the SQL Server environment.
3. Expand the **Groups** heading in the array's pane window and click the **Logical Units** link.

Notice that the right pane has three tabs: **Logical Units**, **RAID Groups** and **DP Pools**.

- Click the **DP Pools** tab and click the **Create Pool** button.



- Select a **RAID Level** and **Combination**.

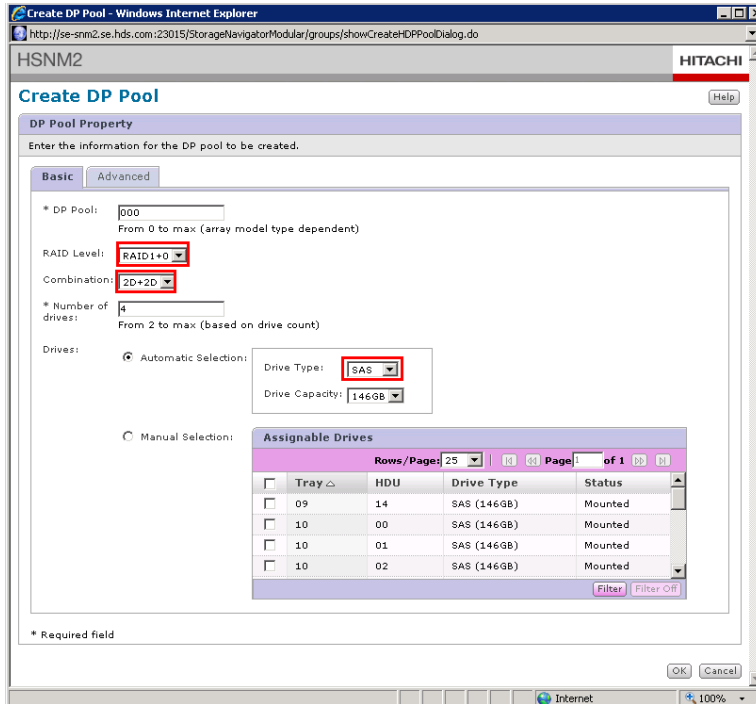
The number of drives automatically changes.

Note that after a RAID level and combination are selected, all future RAID groups added to the Dynamic Provisioning pool also have the same settings.

- Select the **Automatic Selection** radio button.

If you have different types of drives installed in the storage system (for example SATA and SAS, or different capacities) be sure to select the correct type so that Storage Navigator Modular 2 software can select the right type of hard drive.

Enabling automatic selection means that Storage Navigator Modular 2 software selects the next available drives of the type shown in the box. Hitachi Data Systems recommends enabling this setting.



7. Select the **Advanced** tab and modify any of the settings according to your environment requirements and click **OK**.

A message displays indicating that the pool was successfully created.

8. Click the **Close** button.

Note that the pool immediately starts the formatting process in the background.

9. To increase the pool capacity, click the **DP Pools** tab, select the pool's checkbox and click the **Add Pool Capacity** button.
10. Select the **Automatic Selection** radio button and click **OK**.

Enabling automatic selection means that Storage Navigator Modular 2 software selects the next available drives of the type shown in the box. Hitachi Data Systems recommends using this setting.

Dynamic Provisioning pool capacity increases and is made available. Be aware that the pool might momentarily show a Normal (Formatting (0%)) status while it formats the newly created portion of the pool. Nevertheless, any previously data written to the pool is not affected and the new storage capacity is already available to the Dynamic Provisioning pool.

Repeat these steps to add capacity to the Dynamic Provisioning pool.

Creating a Dynamic Provisioning LU

Follow these steps to create a Dynamic Provisioning LU:

1. Expand the **Groups** heading in the storage system's pane window and click the **Logical Units** link.
2. Click the **DP Pools** tab and click the pool number to view the pool properties.
3. Click the **Create LU** button.
4. Enter the required LU size and click **OK**.

This Dynamic Provisioning LU can now be assigned to the host group the same way as a standard LU.

The screenshot displays the Hitachi Storage Navigator Modular 2 web interface. The left sidebar shows a tree view with 'Logical Units' selected. The main content area shows the 'DP Pool-000' properties. Below the properties is a table for 'Logical Units' with columns for LUN, Capacity, Consumed Capacity, Stripe Size, Cache Partition, Pair Cache Partition, and Status. The table is currently empty, showing 'No Object'.

DP Pool	000	Capacity	264,000	Total Consumed	0,000
RAID Level	RAID1+0	DP Pool Consumed Capacity Alert	40%	Early Alert	40%
Combination	2D+2D	Depletion Alert	50%	Warning	100%
Page Size	32MB	Over Provisioning Threshold	130%	Limit	130%
Drive Type	SAS			Notification	No
Status	Normal(Formatting(38%))				

LUN	Capacity	Consumed Capacity	Stripe Size	Cache Partition	Pair Cache Partition	Status
No Object						

Virtual Server Installation

VMware vCenter allows for easy installation and duplication of virtual machines. When using a building block approach, VMware recommends creating a single virtual machine and utilizing it for duplication to simplify deployment. This section assumes that vCenter, vSphere client and ESX hosts are installed and all ESX hosts are added to vCenter.

Creating a Virtual Machine Template

Because multiple virtual machines of identical configuration are required for this solution, using a virtual machine template eases deployment. For more information about virtual machine templates, see VMware's support [Web site](#).

Follow these high-level steps to create a virtual machine template using vCenter:

1. Use the vSphere client to log in to the appropriate vCenter console.
2. Select the appropriate ESX host and create and assign a LU to it.
3. Create a datastore for the LU.
4. Create a virtual machine and install Windows Server 2008.
5. Add all patches and any additional software needed for the VM guest operating system.
6. Install VMtools on the VM.
7. Shut down the VM, right click the VM in vCenter and select the **Convert to Template** option in the **Template** menu.

Additional virtual machines can be deployed using this template. Be sure to use the customization wizard to generate a new SID.

Deploying SQL Server Virtual Machines

This section focuses on implementing SQL Server 2008 on virtual machines. These are high level steps that assume the user has a good understanding of how to execute tasks on VMware's vCenter.

Table 5 describes an environment built using the building block architecture that supports the SQL Server instance. A 500GB building block architecture is used in this solution. The database uses a scalable approach from a virtual machine, storage performance and capacity perspective.

Table 5. Building Bock Architecture Supporting a 500GB Database Environment Using Standard LUs

<i>Resource</i>	<i>Details</i>
Virtual machine	32GB memory 2 vCPU
Storage system - Hitachi Adaptable Modular Storage 2300	16 x 146GB 15K RPM SAS disks
RAID groups	1 x RAID-1+0, 4D+1P 1 x RAID-1+0, 6D+6D 1 x RAID-1+0, 2D+2D
LUs	1 x 100GB on RAID5-4D+1P - OS 1 x 625GB on RAID-1+0 6D+6D - database 1 x 62.5GB on RAID-1+0 6D+6D - tempdb 1 x 160GB on RAID-1+0 2D+2D - T-logs

OLTP workload conducted in Hitachi Data Systems' lab shows that two vCPU and 32GB of memory on the virtual machine adequately supports this configuration and also contains some excess CPU and memory headroom for burst activity depending on your user workload.

Deploying SQL Server 2008 on Virtual Machines

Follow these steps to deploy SQL Server instances on a VM in vCenter:

1. Use the vSphere client to log in to the appropriate vCenter console.
2. Select the appropriate ESX host and deploy a VM using the template created earlier.
3. Use the Guest Customization wizard to customize host name, domain name, IP address and other networking parameters.
4. Turn the VM on.
5. Install SQL Server 2008 along with appropriate patches.
6. Add LUs from the 2000 family storage system to the ESX host.

The newly added storage is discovered by the ESX host.

For more information about adding LUs to the ESX host, see the Configuring the Adaptable Modular Storage 2000 section in this paper.

7. Follow these steps to create datastores on all LUs:
 - a) Select an ESX host and navigate to the **Configuration** tab.
 - b) Click **Add Storage** and use the wizard to add the datastore.

The datastores are created. In this example, the LU for the database file is 625GB, therefore a block size of 4MB is required to support this vDisk size.

For more information, see [VMware's Fibre Channel SAN Configuration Guide](#).

- c) Follow these steps to select the round robin path policy algorithm:
 - i. Right-click the datastore and select **Properties**.
 - ii. Click **Manage Paths**.
 - iii. In the drop-down menu, select **Round-Robin (VMware)**
8. Follow these steps to assign vDisks to the virtual machine:
 - a) Select the desired virtual machine and edit the virtual machine settings.
 - b) Select **Add**.
 - c) Select **Hard Disk** and click **Next**.
 - d) Create a new Virtual Disk and click **Next**.
 - e) Select **Specify a datastore** and browse to the datastore for a storage group.
 - f) Enter the size as the maximum capacity of the datastore and click **Next**.
 - g) Accept the advanced options and click **Next**.
 - h) Click **Finish**.
 - i) Repeat steps a through h for all vDisks
 - j) Click **OK**.

For more information, see the vSphere 4 [product support page](#).

After all the vDisks are assigned to the VM, Windows discovers the new disks. Use standard Windows practices to configure these disks using the Disk Management MMC. The disks can be either mounted using a drive letter or on a mount point. Because SQL databases can scale to require a number of LUs for a given SQL Server instance, Hitachi Data Systems recommends using mount points due to the drive letter limitation. For more information about standard Windows practices, see Windows online help.

Validating Your Deployment

Like any other type of SQL Server deployment, test your environment to ensure that the configuration meets your performance requirements.

Microsoft offers the SQLIO and the SQLIOSim tools to help validate your virtualized SQL Server environment. Nevertheless, it is important that you test your own environment in a close to production mode. While SQLIO and SQLIOSim provide a generic view of what the environment is capable of, your specific database workload and user load can greatly differ from SQLIO. The Iometer tool is also a good option when you have deep knowledge of your specific database workload. For more information about utilizing tools to validate your SQL Server environment, visit the [Deploying SQL Server 2005 with SAN #3](#) article from SQLCAT's Web site.

To download these tools, visit their specific sites.

- **SQLIO** — [SQLIO Disk Subsystem Benchmark Tool](#)
- **SQLSim** — [How to use the SQLIOSim utility to simulate SQL Server activity on a disk subsystem](#)
- **Iometer** — [SourceForge](#)

Always start by deploying the environment as a test or development environment and test the SQL Server instance to ensure the correct levels of IOPs are achieved within acceptable latency ranges. Keep in mind that a holistic view of your environment is necessary to appropriately identify any possible bottlenecks.

Monitoring and Managing the Solution

Tools from Hitachi, Microsoft, and VMware can ease management, reduce complexity and provide peace of mind that your solution was planned and deployed successfully.

Managing a Growing Environment

Most SQL Server environments eventually need to add storage groups due to the increase in size of the databases or need to increase the available amount of drives for better IOPS. This solution uses a building block approach to simplify scaling your environment. You can grow your virtual SQL Server instances by adding additional RAID groups and LUs to the environment or simply adding more Dynamic Provisioning LUs when the deployment uses Hitachi Dynamic Provisioning software.

- **Database capacity** — Increase the capacity of your databases by adding additional RAID groups and LUs, or by simply adding more Dynamic Provisioning LUs to your Dynamic Provisioning pool.
- **Virtual machines** — Increase the number of SQL Server instances on your environment according to your hardware needs.
- **ESX hosts** — Each ESX host can accommodate up to 4 virtual machines of varying capacities and workload profiles.

For this environment, the ESX host contained 128GB of memory and four quad-core AMD Opteron CPUs. When planning to scale your environment, it's also important to consider your current size and the database workload expected from users to be added when deciding which increment is best. In all cases, additional RAID groups and LUs must be provisioned on the 2300.

VMware Product Features

VMware provides product features that can help manage the entire infrastructure:

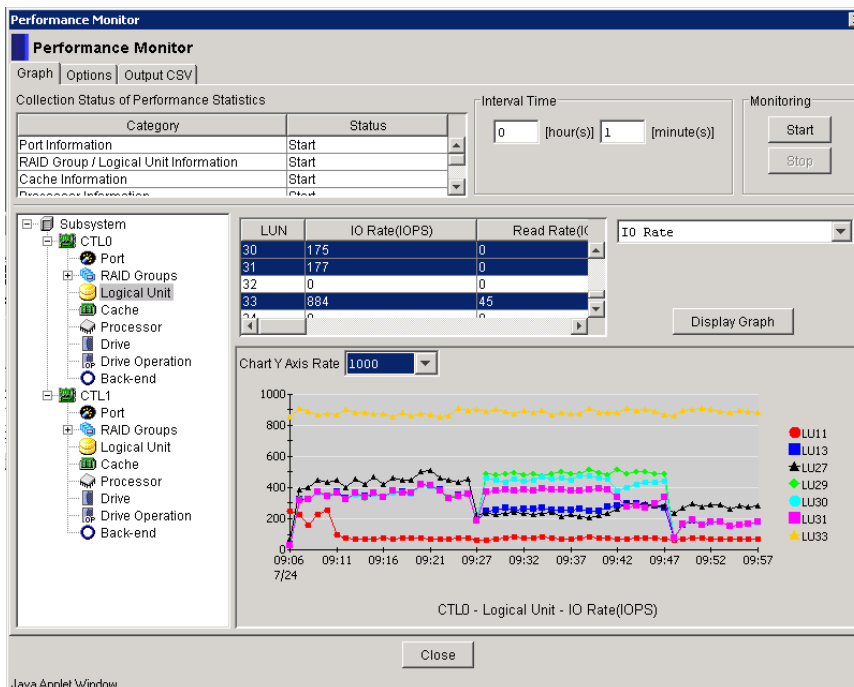
- **vMotion** — Allows for non-disruptive migration of both virtual machines and storage. Its performance graphs allow for monitoring of resources, virtual machines, resource pools and server utilization.
- **Distributed Resource Scheduler (DRS)** — Monitors resource utilization and intelligently allocates system resources as needed.
- **High Availability (HA)** — Monitors for hardware and OS failures and automatically restarts the virtual machine, providing cost effective failovers.
- **Fault Tolerance (FT)** — Provides continuous availability for applications by creating a live shadow instance of the virtual machine that stays in lockstep with the primary instance. If there is a hardware failure the shadow instance instantly takes over and eliminates even the smallest data loss.

Hitachi Performance Monitor Feature

Hitachi Performance Monitor feature is part of the Storage Navigator Modular 2 software package. It acquires information on the performance of RAID groups, LUs and other elements of the storage system while tracking the utilization rates of resources such as hard disk drives and processors. Information is displayed with line graphs in the Performance Monitor feature windows or saved in .csv files. Use Performance Monitor feature to validate that the underlying storage system is not taxed by the SQL environment.

Figure 5 shows the monitoring of an environment using Performance Monitor feature's GUI.

Figure 5. Performance Monitor Feature User Interface

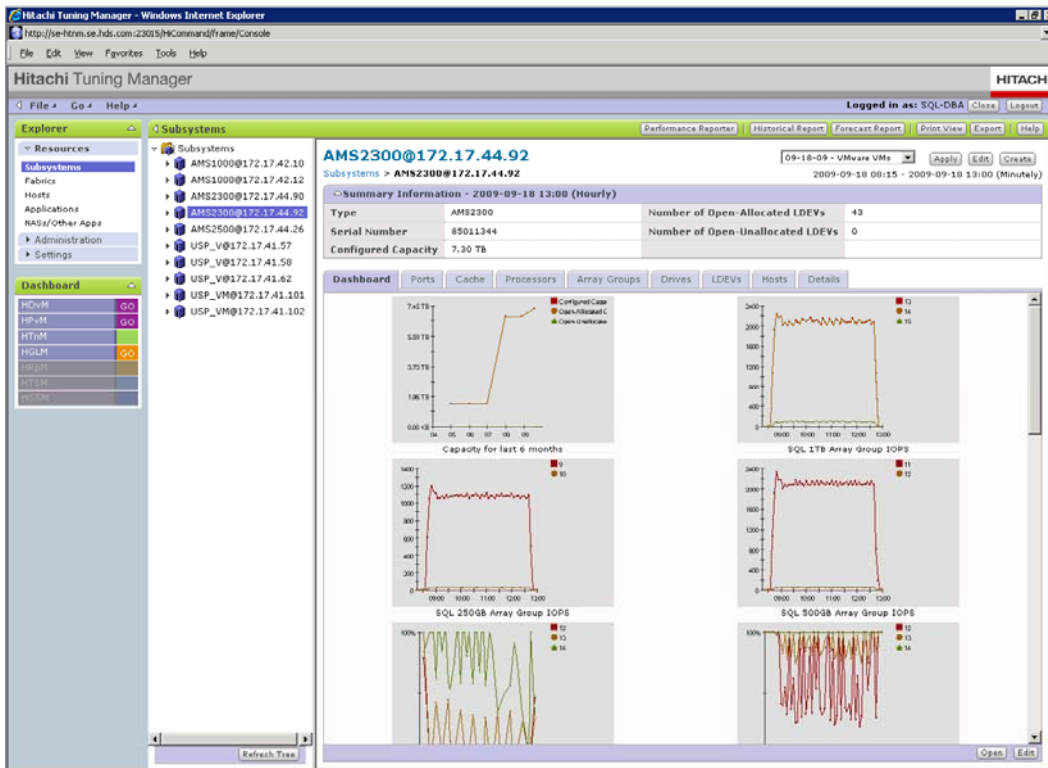


Hitachi Tuning Manager Software

Hitachi Tuning Manager software enables you to proactively monitor, manage and plan the performance and capacity for the Hitachi modular storage that is attached to your vSphere servers. Hitachi Tuning Manager software consolidates statistical performance data from the entire storage path. It collects performance and capacity data from the operating system, SQL Server instances, switch ports, and storage system information such as RAID groups and LUs IOPS and latency measures. Hitachi Tuning Manager software provides you with a complete performance picture of your environment. It provides historical, current and forecast views of these metrics. For more information about Hitachi Tuning Manager software, see the [Hitachi Data Systems support portal](http://www.hitachidts.com/support).

Figure 6 shows the monitoring of a SQL environment using Hitachi Tuning Manager software.

Figure 6. Hitachi Tuning Manager Software User Storage Dashboard





Corporate Headquarters 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Asia Pacific and Americas 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Europe Headquarters Sefton Park, Stoke Poges, Buckinghamshire SL2 4HD United Kingdom
Contact Information: + 44 (0) 1753 618000 www.hds.com / info.uk@hds.com

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

All other trademarks, service marks and company names mentioned in this document or Web site are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems. This document describes some capabilities that are conditioned on a maintenance contract with Hitachi Data Systems being in effect and that may be configuration dependent, and features that may not be currently available. Contact your local Hitachi Data Systems sales office for information on feature and product availability.

© Hitachi Data Systems Corporation 2010. All Rights Reserved.
AS-023-01 March 2010