

# Virtualizing Microsoft® SQL Server 2008 on the Hitachi Adaptable Modular Storage 2000 Family Using Microsoft Hyper-V

Implementation Guide

*By Eduardo Freitas and Ryan Sokolowski*

February 2010



## Summary

Deploying a virtualized Microsoft® SQL Server 2008 environment with Microsoft Hyper-V 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 Hyper-V 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](mailto:SolutionLab@hds.com). Be sure to include the title of this white paper in your e-mail message.



# Table of Contents

<b>Solution Overview .....</b>	<b>1</b>
<b>Hardware Resources.....</b>	<b>4</b>
Hitachi Adaptable Modular Storage 2000 Family .....	4
Servers.....	4
Storage Area Network.....	4
<b>Software Resources.....</b>	<b>6</b>
<b>Planning for Deployment.....</b>	<b>7</b>
Deployment Scenarios .....	7
Sizing the Hyper-V Servers.....	7
Sizing SQL Server 2008 Virtual Machines .....	7
Planning Storage Performance and Capacity .....	8
<b>Deploying the Solution .....</b>	<b>9</b>
Configuring the SAN .....	10
Storage Configuration .....	10
Using Hitachi Dynamic Provisioning Software.....	17
Virtual Server and Virtual Machine Configuration.....	20
Validating Your Deployment.....	25
<b>Monitoring and Managing the Solution .....</b>	<b>26</b>
Managing a Growing Environment.....	26
Hitachi Performance Monitor Feature .....	26
Hitachi Tuning Manager Software.....	27
Hyper-V Protection Strategies.....	28

# Virtualizing Microsoft® SQL Server on the Hitachi Adaptable Modular Storage 2000 Family Using Microsoft Hyper-V

## Implementation Guide

Deploying a virtualized Microsoft® SQL Server 2008 environment with Microsoft Hyper-V 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 Hyper-V 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 Hyper-V 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 Hitachi Storage Navigator Modular 2 software, Microsoft Hyper-V, 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

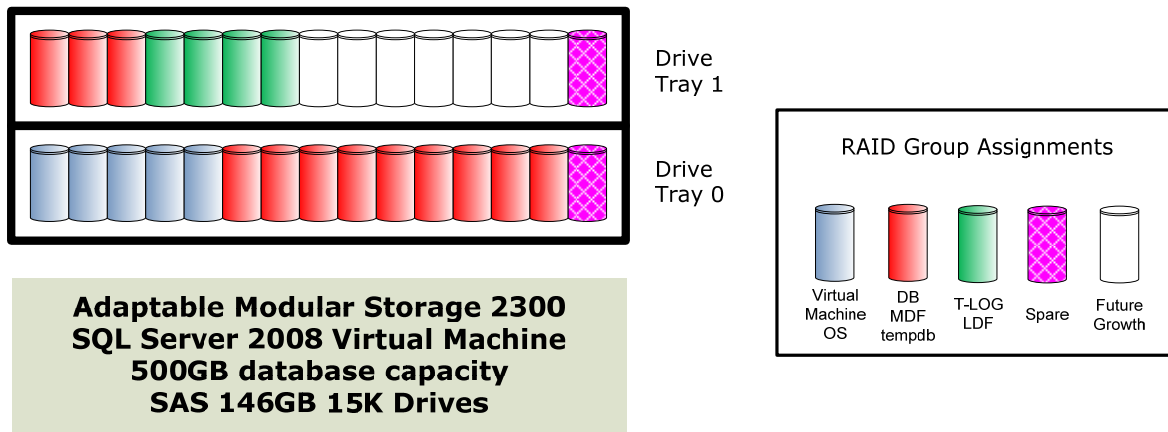
The building block architectures described in this white paper were developed with ease of management and scalability in mind. The architectures consist of a Hyper-V virtual machine (VM) running Microsoft Windows Server 2008 and Microsoft SQL Server 2008 with Service Pack 1 with storage hosted on the Hitachi Adaptable Modular Storage 2300. Logical units (LUs) from the 2300 are allocated to the Windows Server 2008 host on which the Hyper-V role is installed. The LUs are presented to the SQL Server VMs as the pass-through disks and can then be partitioned and used in file systems for the SQL Server user databases, tempdb, 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 Microsoft's 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. For more

information about Microsoft guidelines, see Microsoft's [Troubleshooting Performance Problems in SQL Server 2008](#) white paper. 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.

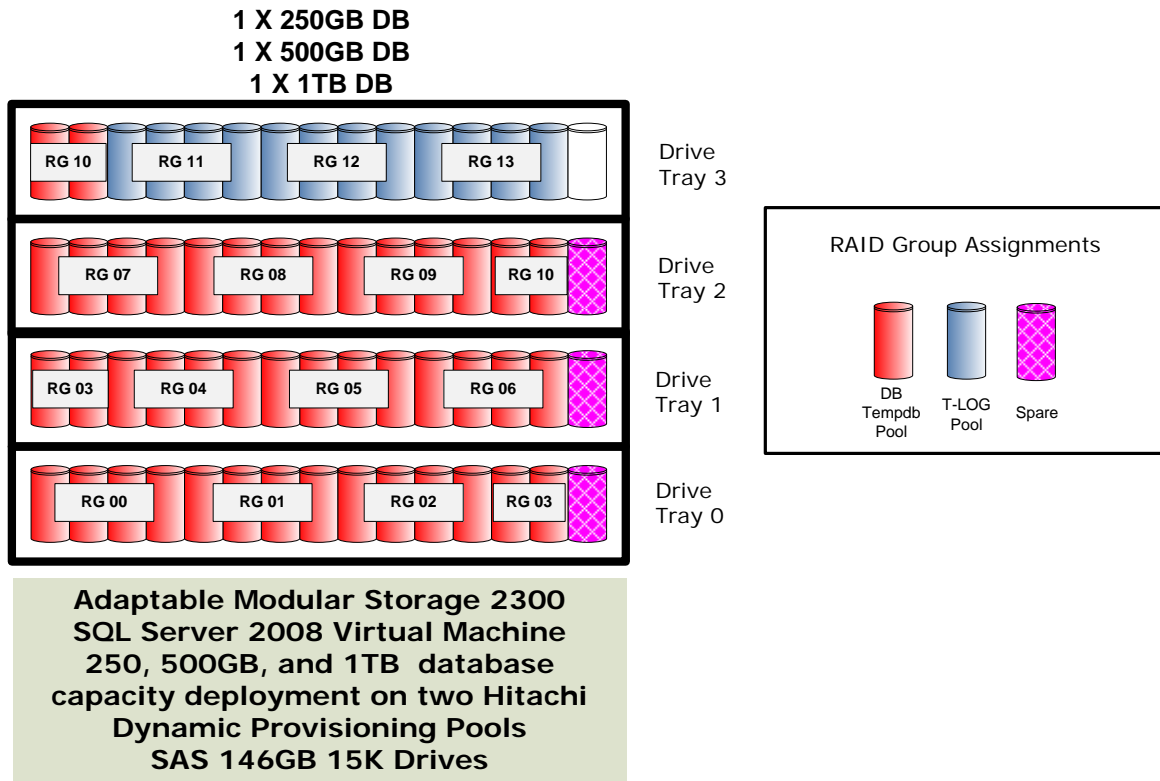
Figure 1 shows the underlying storage configurations for the 500GB database building block for the traditional storage configuration that does not use the dynamic provisioning approach.

**Figure 1. Standard Provisioned 500GB Storage Building Block**



For storage configurations using Hitachi Dynamic Provisioning software, use 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 with 250GB, 500GB, and 1TB capacities, use a total of 14 RAID-1+0 (2+2) groups, 11 of which are used for the database and tempdb volumes, and three 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 in Figure 1 with regards to performance, capacity and scalability requirements. 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



Note that the RAID group used for the virtual machine OS is not shown in this picture.

For more information about the building block architecture and Hitachi Dynamic Provisioning software, see the [Virtualizing Microsoft® SQL Server on the Hitachi Adaptable Modular Storage 2000 Family Using Hyper-V Reference Architecture](#) 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 used for your virtualized SQL Server deployment. Determining which member of the 2000 family is best for your environment requires evaluating your overall capacity and connectivity requirements.

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

### Servers

Use two enterprise class level servers when deploying this virtualized SQL Server solution in high availability environments. The servers must have enough processor and memory resources to operate the SQL Server environments and meet Windows Server 2008 operating requirements. Hitachi Data Systems and Microsoft recommend that both servers being utilized have the same hardware architecture and configuration and processor and memory resources to accommodate a failover situation.

For more information about Hyper-V server hardware requirements, see the [Microsoft Hyper-V Server 2008 R2 – System Requirements](#) Web site.

### Storage Area Network

For this solution, Hitachi Data Systems recommends connecting the Hyper-V 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 multiple blades that can support high availability and redundancy.

In addition, Hitachi Data Systems recommends configuring a minimum of two redundant paths from the Hyper-V host to the 2300. Each Hyper-V host requires at least two host bus adapters (HBAs) for high availability purposes. Hitachi Data System Dynamic Link Manager'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. Alternatively, you can use the round-robin algorithm available with the Windows' MPIO feature. Hitachi Data Systems lab tests showed that using either MPIO's round-robin policy or Hitachi Dynamic Link Manager software's round robin algorithm proved to be the best choice for SQL Server OLTP workloads compared to other available policies.

Figure 3 illustrates the storage area network configuration.

**Figure 3. Minimum Storage Area Network Configuration**

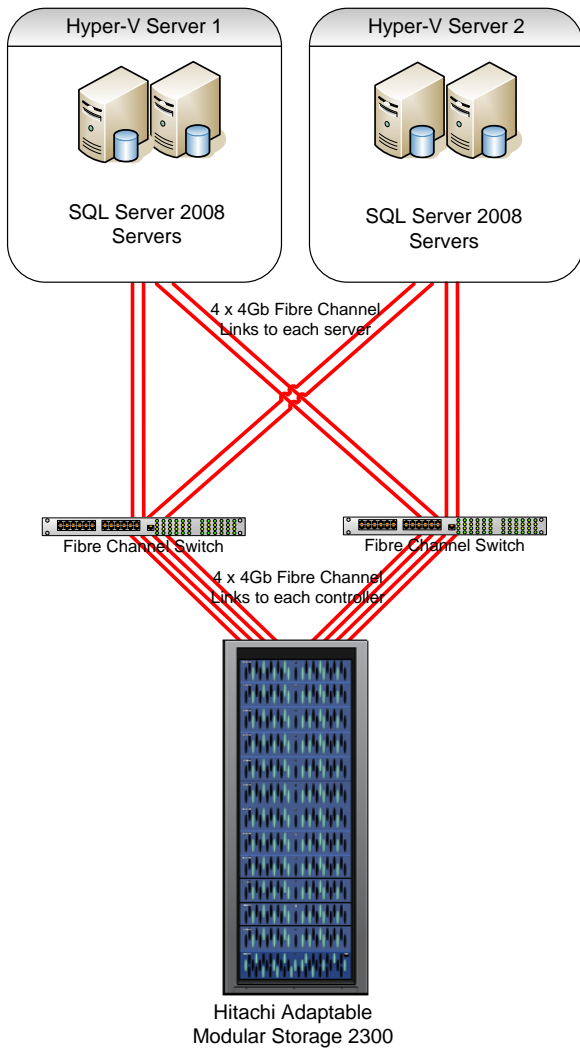


Table 1 lists the hardware used in testing this architecture.

**Table 1. Tested Deployment Hardware Resources**

<i>Hardware</i>	<i>Configuration</i>
Hitachi Adaptable Modular Storage 2300 storage system	Microcode 0883A-M* 2 controllers 4 disk trays 15 146GB 15K RPM SAS disks per tray 8GB 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 Hitachi supported levels of microcode and FOS levels

\*\*Use the latest available Hitachi and Microsoft 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>
Windows Server 2008 (for Hyper-V server)	Datacenter edition, R2
Windows Server 2008 (for SQL Server virtual machines)	Enterprise edition, Service Pack 2
SQL Server 2008	Enterprise edition, Service Pack 1
Hitachi Storage Navigator Modular 2	7.0*
Hitachi Dynamic Provisioning license	Microcode dependent
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.

\*\*Note: While Hitachi Device Manager software and Hitachi Tuning Manager software are not required software resources for deployment of this solution, they provide features that enable you to more effectively manage the storage system and monitor the performance of your environment.

For more information about the software used to deploy this solution, see the [Virtualizing Microsoft® SQL Server 2008 on the Hitachi Adaptable Modular Storage 2000 Family Using Hyper-V Reference Architecture Guide](#) white paper 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 Microsoft Hyper-V 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 Microsoft Hyper-V Server. For more information about SQL Server consolidation, see the Microsoft white paper, [SQL Server Consolidation Guidance](#).

Backup and restore (or recovery) schemes 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 focuses on deploying SQL Server 2008 on Microsoft Hyper-V Server on a 500GB capacity building block. While this building block approach supports different deployments, the information in this white paper enables you to deploy a single 500GB building block. Larger environments can be built by repeating the steps described in this guide with regards to capacity.

## Sizing the Hyper-V Servers

To establish the hardware requirements for your Hyper-V Server, you must first determine how many virtual machines will run on each physical Hyper-V server. You must also determine the virtual CPU, memory and capacity requirements for each of the virtual machines you are planning to deploy. You also need to consider whether the Hyper-V Server will be part of a Hyper-V failover cluster and whether features like Hitachi Storage Cluster with quick or live migration will be implemented for High Availability. Note that at this point, Hyper-V is limited to 64 VMs per clustered node and 384 VMs per host in standalone environments.

For more information, see the [Virtualizing Microsoft® SQL Server 2008 on the Hitachi Adaptable Modular Storage 2000 Family Using Hyper-V Reference Architecture Guide](#) and [Optimizing High Availability and Disaster Recovery with Hitachi Storage Cluster in Microsoft® Virtualized Environments with Live Migration](#) white papers.

### *Virtual CPUs and Memory*

Virtual processors are the processors within the virtual machines, and they equate to one physical processor core. The Hyper-V hypervisor is responsible for virtualizing the processors for a given virtual machine. Microsoft recommends ensuring that the total number of virtual processors assigned to all the virtual machines is equal to or less than the total number of cores on the physical Hyper-V server. For example, if you are using a server that has two quad-core CPUs (a total of eight cores), and you need to deploy four virtual machines on the server, each machine should have at most two virtual processors assigned to it. While it is possible to over-commit processors to VMs, best practice is to not over-commit unless the host resources are regularly monitored. Over-committing can cause a performance bottleneck to the VMs.

Unlike the processors, memory cannot be over-committed in a Hyper-V environment. By not allowing the memory to be over-committed to virtual machines, Hyper-V ensures that no performance bottleneck occurs in your virtual environment.

## 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 virtual processors and the amount of memory required. You must ensure that the processor and memory requirements for the SQL Server virtual machines match the recommended configurations as per Microsoft requirements. Network connectivity also plays an important part of your SQL Server virtual environment, and must be allocated using virtual switches. For more information about the hardware and software requirements for installing SQL Server 2008, see the [Hardware and Software Requirements for Installing SQL Server 2008](#) Books Online Web site. Also keep in mind that different SQL Server 2008 editions support different maximum numbers of processors and that a maximum number of logical processors (four) can be assigned to a given virtual machine. For more information, see the [Maximum Number of Processors Supported by the Editions of](#)

[SQL Server Books Online Web site](#). Microsoft recommends deploying a single instance of SQL Server per virtual machine. Deploy additional virtual machines if you need more SQL Server instances to support your environment for improved performance and security isolation.

### *SQL Server Processor and Memory Considerations*

Establishing the number of processors and amount of memory a given SQL Server needs depends on many factors, including the SQL Server 2008 edition, the backup strategy, the utilization of data compression, reporting services and, most importantly, user workload for the instance user 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 virtual processors 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 virtual processors assigned to the virtual machine as per Microsoft's best practice. Keep in mind that unlike a physical processor, which contains many cores, a single virtual processor only provides a single logical processor core to a virtual machine.

For example, for this solution's deployed architecture, two virtual processors are allocated to the SQL Server virtual machine, which means that 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 tempdb file location and the addition of one more tempdb data file can be made by executing the following SQL statements:

```
use tempdb
go

-- 2 virtual processors --

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 LU that is assigned to the Hyper-V server 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 Web site.

## 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 (6D+6D) group that provides 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 use a RAID-1+0 (2D+2D) group that also provides for a 20 percent growth space. For more information, see the Virtualizing SQL Server on the Hitachi Adaptable Modular Storage 2000 Family Using Hyper-V 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. The building block described in this guide follows Hitachi Data Systems and

Microsoft best practices for planning storage capacity for SQL Server databases and logs. Consider the following when sizing the storage and virtual machine for performance and capacity:

- Overall capacity needed for data files and transaction logs including planned growth for the SQL Server instances.
- IOPS needed for SQL Server data files and transaction logs.
- Set tempdb file growth a reasonable value as per Microsoft's best practices. Setting the file growth to small values can negatively affect the performance of your environment. For more information, see the [Optimizing tempdb performance](#) Books Online Web site.
- SQL Server instance host CPU and memory requirements for the expected database workloads.
- Network bandwidth requirements each SQL Server virtual machine.

Be sure to account for overhead when sizing the databases and logs. A typical overhead measure is to 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 Hyper-V Server. Use global unique identifiers (GUIDs) whenever possible. This allows for greater flexibility in that you are not limited to 26 drive letters on the Hyper-V server.

## 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 increase the capacity of a given SQL instance or to deploy 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 virtual processors
Storage systems	1 Hitachi Adaptable Modular Storage 2300 23 146GB 15K RPM SAS disks Note: 2 disks are used as hot spares
RAID groups	1 RAID-1+0 (6D+6D) 1 RAID-1+0 (2D+2D) 1 RAID-5 (4D+1P)
LUs	1 625GB on RAID-1+0 (6D+6D) for database files 1 62.5GB on RAID-1+0 (6D+6D) for tempdb files 1 160GB on RAID-1+0 (2D+2D) for transaction logs 1 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 Hyper-V server.

## Configuring the SAN

This solution requires two or more Hyper-V Servers to form a highly available Hyper-V failover cluster using Hitachi Storage Cluster with Microsoft quick or live migration. Configuring the SAN in a Hyper-V environment requires careful consideration of all Hyper-V servers in the cluster. Configure the SAN zones so that each of the four ports on the Hyper-V server running the SQL Server instances has visibility to one of the four ports on the 2000 family storage system: two ports from each storage system's controller. See Figure 3 for an illustration of the connection between the Hyper-V servers and the storage system.

Note that while four dedicated ports from the host and four from the storage system were used 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 exceeding the IOPS capability of any port and therefore not creating a bottleneck in the environment. In this case, zone each server port to a port in the storage system, each in a different controller for high availability purposes.

Ensure that both the storage and host Fibre Channel ports are configured for point-to-point connections. 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 Hyper-V server and the storage system ports.

**Table 4. Path Configuration for a Single Hyper-V Server**

<i>Hyper-V Server</i>	<i>Host HBA Number</i>	<i>Switch</i>	<i>Zone Name</i>	<i>Storage System Port</i>	<i>Storage System Host Group</i>
Hyper-V Server 1	HBA 1 port 1	Brocade 5100 - 1	HyperV_1_HBA1_1_AMS2300_0A	0A	HyperV_1_HBA1_1
Hyper-V Server 1	HBA 1 port 2	Brocade 5100 - 2	HyperV_1_HBA1_2_AMS2300_1A	1A	HyperV_1_HBA1_2
Hyper-V Server 1	HBA 2 port 1	Brocade 5100 - 1	HyperV_1_HBA2_1_AMS2300_0B	0B	HyperV_1_HBA2_1
Hyper-V Server 1	HBA 2 port 2	Brocade 5100 - 2	HyperV_1_HBA2_2_AMS2300_1B	1B	HyperV_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 software. 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 storage for the SQL Server; however, the same procedure can be used when deploying the storage for the virtual machines' operating system.

## 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 in your data center.  
Ensure that you have Modify permissions within 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 the **Logical Units** link.
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.

The screenshot shows the 'Create RAID Group' dialog box in a web browser. The 'RAID Level' dropdown is set to 'RAID1+0' and the 'Combination' dropdown is set to '6D+6D'. The 'Number of Parity Groups' is set to 1. The 'Automatic Selection' radio button is selected, and the 'Drive Type' is 'SAS' and 'Drive Capacity' is '146GB'. An 'Assignable Drives' table is visible below.

Tray	HDU	Drive Type	Status
<input type="checkbox"/> 06	05	SAS (146GB)	Mounted
<input type="checkbox"/> 06	06	SAS (146GB)	Mounted
<input type="checkbox"/> 06	07	SAS (146GB)	Mounted

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 a greater amount of 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**.

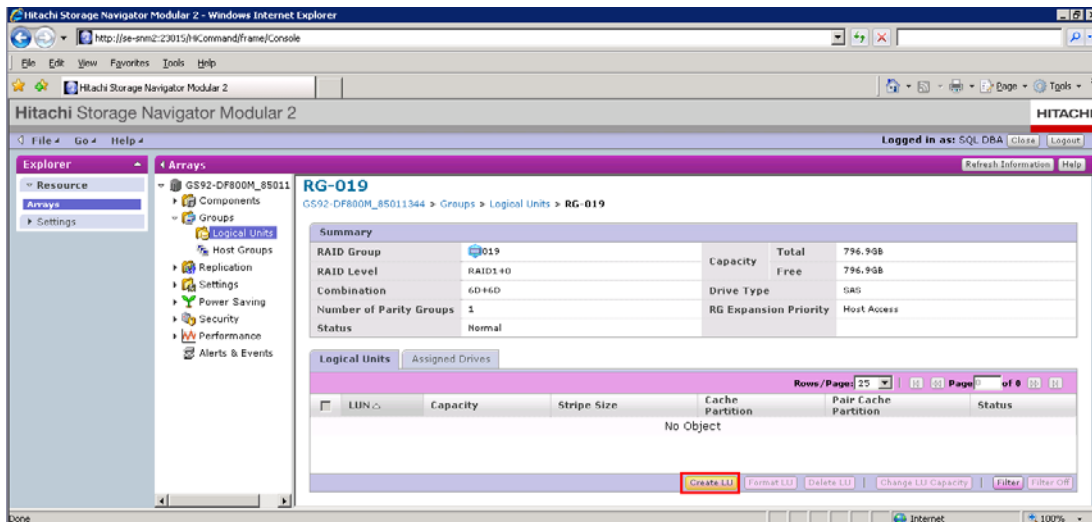
A message displays stating that the RAID group was successfully created.

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



4. Change the **Capacity** information from RG ALL to GB, and enter 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, a minimum of 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

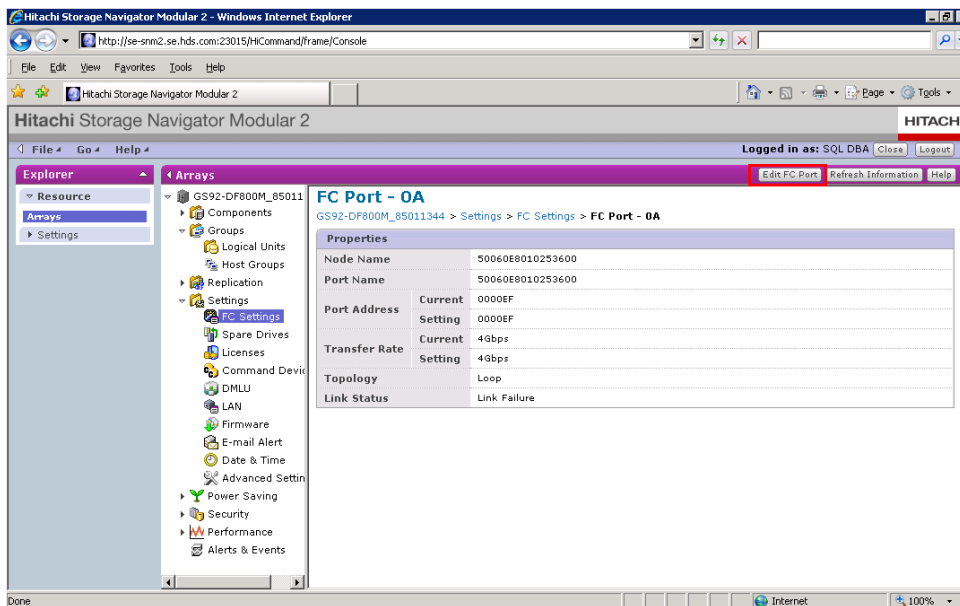
In preparation for this step, all Hyper-V server HBA ports need to be connected to a Fibre Channel switch and appropriately zoned to the storage system ports. Contact your SAN administrator to ensure that the proper zoning configuration is in place in your environment.

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

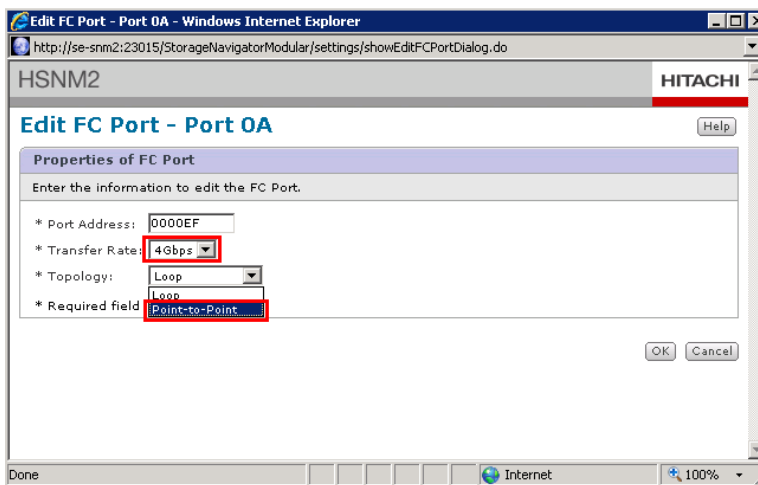
1. Expand the **Settings** heading and click the **FC Settings** link.

Notice the **Topology** column for the ports that were zoned to connect to the Hyper-V server 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 for any host that is connected to the port.

4. Click the **Confirm** button.

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 Hyper-V server's ports before they can connect to the storage system.

After the connection between the storage system and the Hyper-V server is established, the **FC Settings** screen shows all modified ports with a **LinkUp(F\_Port Connected)** status.

### *Enabling Host Group Security*

Enabling host group security enables you to associate specific WWNs to a given port and 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:

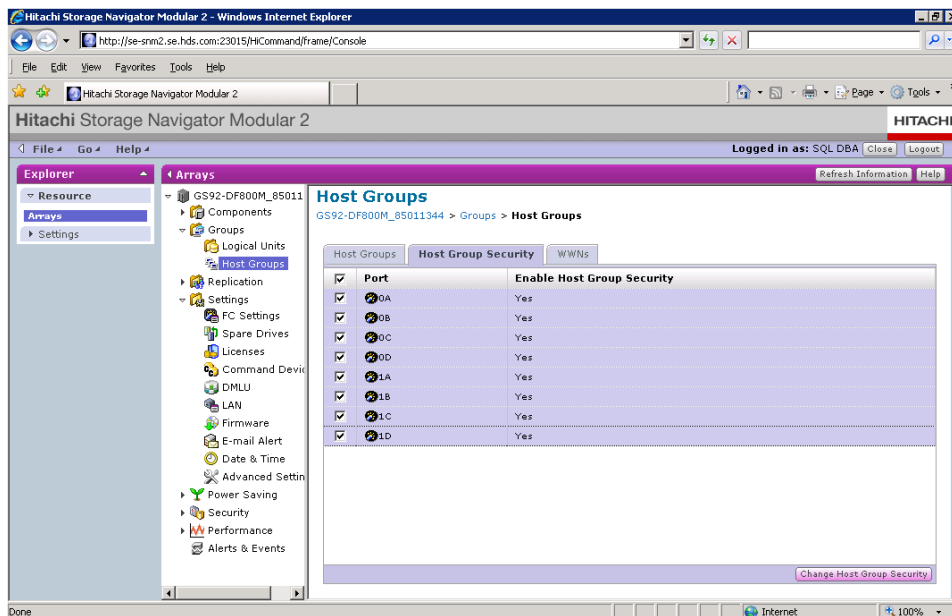
1. Expand the **Groups** heading and click the **Host Groups** link.
2. Click the **Host Group Security** tab and ensure that port security is enabled on all ports.

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** check box and click **OK**.

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

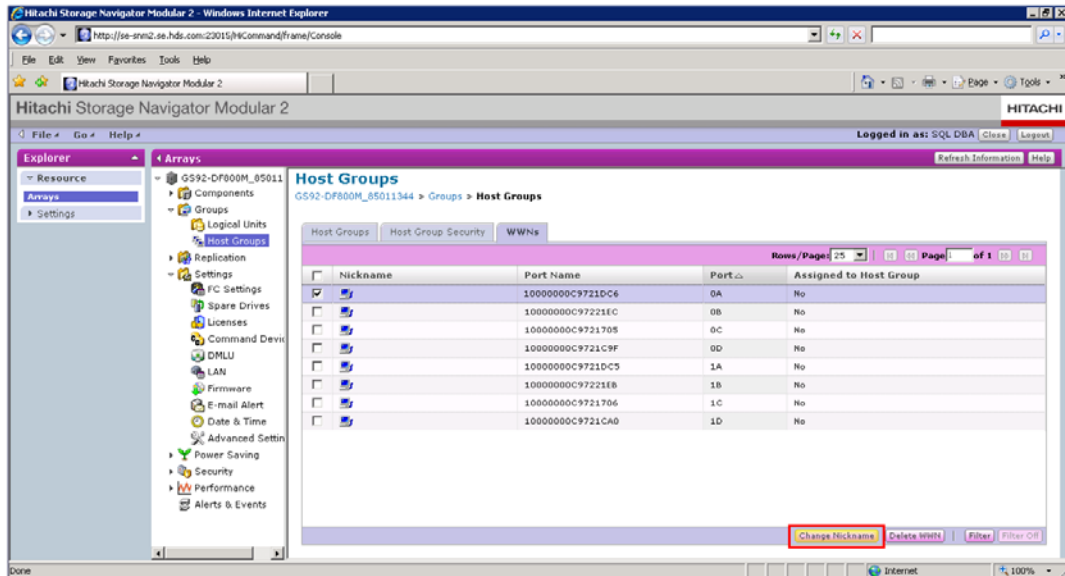


### *Assigning Nicknames to HBA's WWNs*

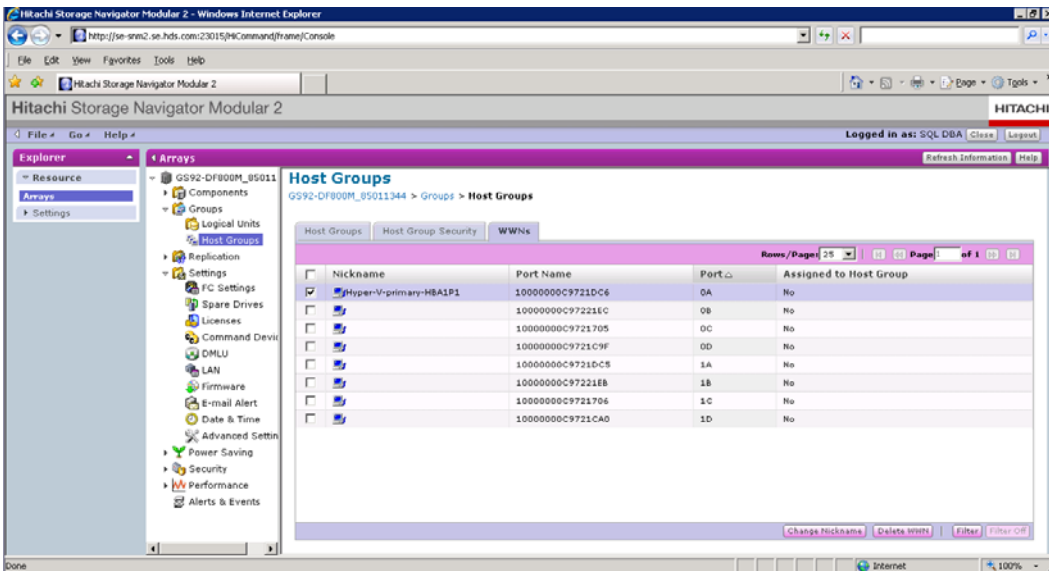
This procedure is optional, however, assigning a nickname to Hyper-V server HBA port's WWNs can facilitate identifying each host and port to ease documenting and troubleshooting your environment. Note that this step assumes that the necessary connections and zones exist on the SAN and that the correct 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.
7. The Storage Navigator Modular 2 software screen automatically refreshes and shows the new nickname for the WWN.

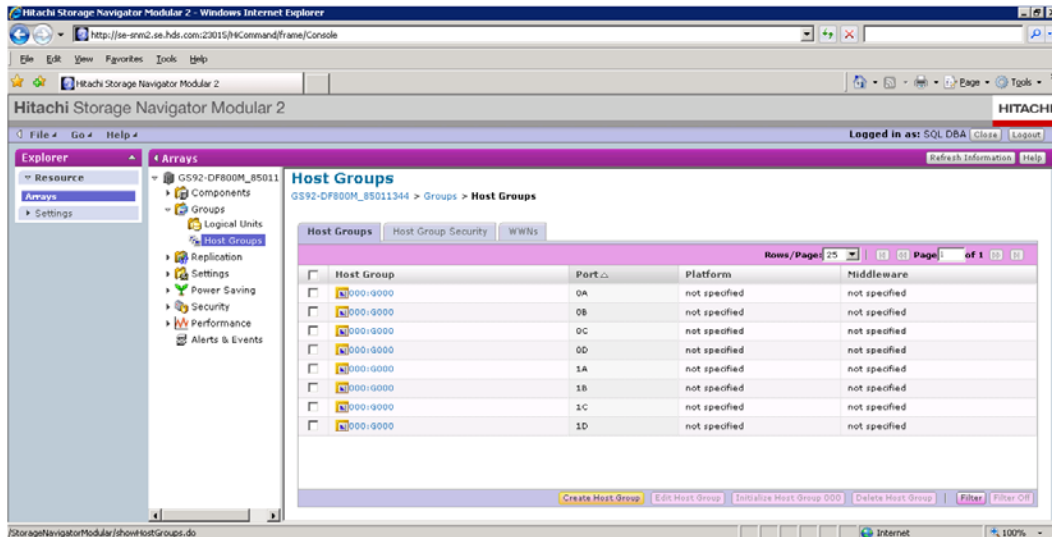


Repeat this step for the Hyper-V Server's remaining WWNs.

## 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.

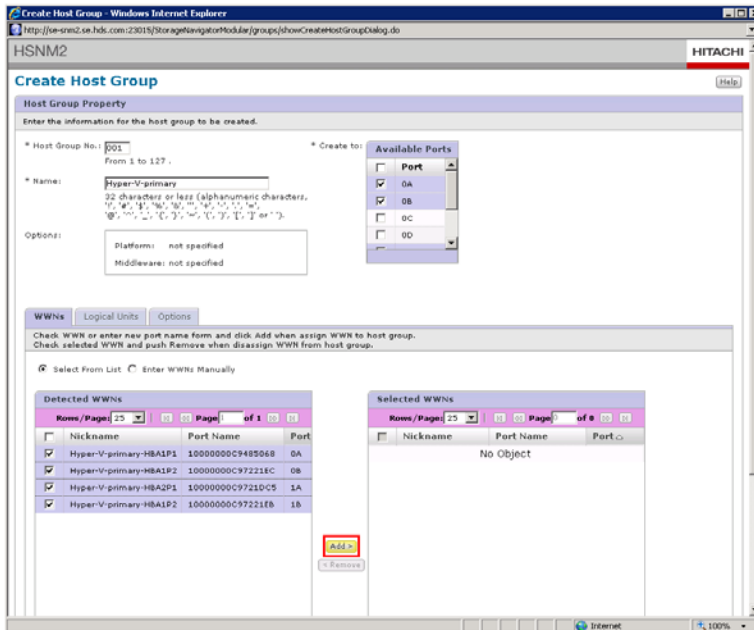
For example, you might use HyperV\_1 for your primary Hyper-V server.

4. Select all of the Fibre Channel ports that are connected to the Hyper-V Server.

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.

5. 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.

6. Click the **Logical Units** tab.
7. 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.

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

Repeat this procedure for other Hyper-V servers 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. For that reason, it is important to use the quick formatting option when creating partitions on a given LU.

Hitachi Storage Navigator Modular 2 software default alert settings for Hitachi Dynamic Provisioning match Hitachi Data Systems recommended minimum 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 reaching an early alert value or 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.

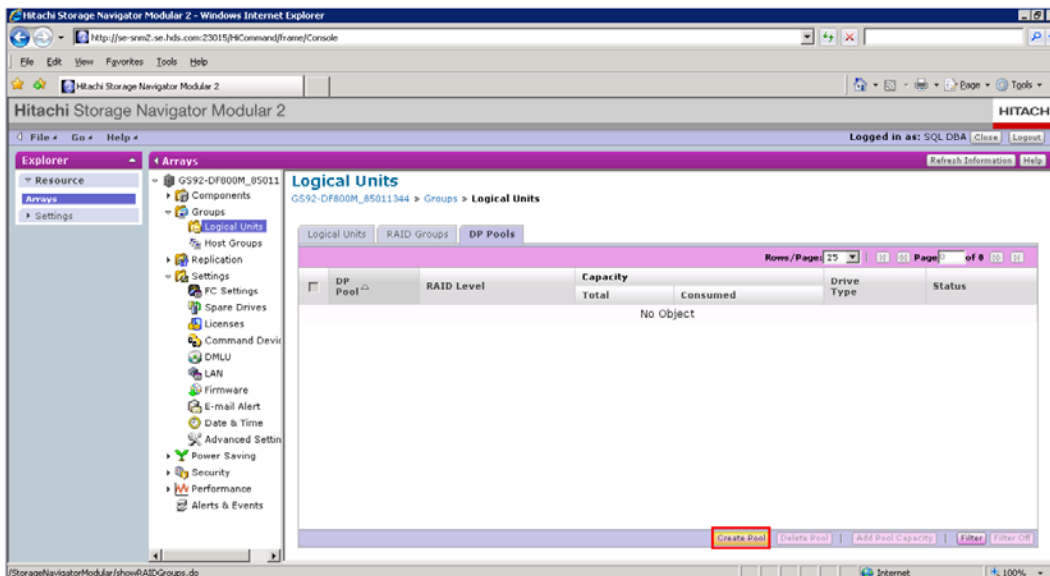
For example, a ratio of 1.4 (140 percent) defines a "comfort limit" ratio by defining more virtual capacity than the current pool capacity can currently support.

The following sections describe 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.
4. Click the **DP Pools** tab and click the **Create Pool** button.



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

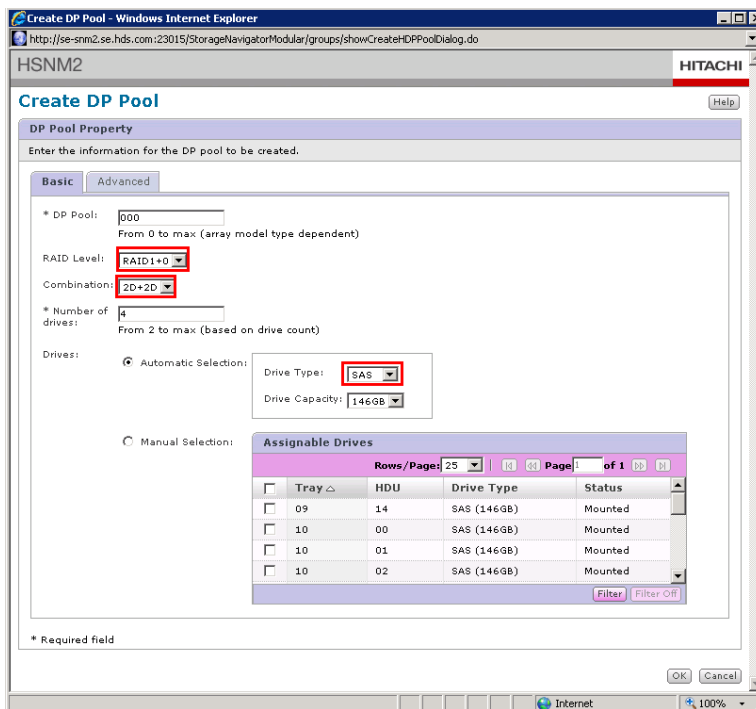
The number of drives automatically changes.

Note that the pool consists of a single RAID group at creation time, and after a RAID level and combination are selected, all future RAID groups added to the Dynamic Provisioning pool also have the same settings.

6. 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 check box 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.

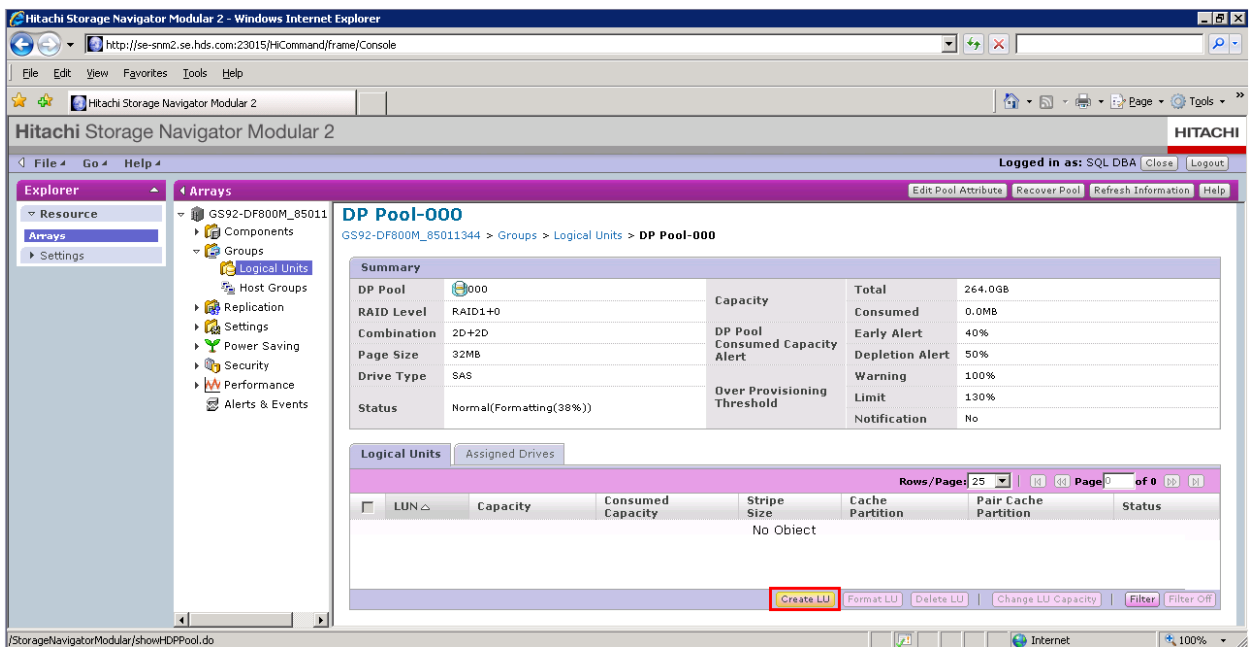
Dynamic Provisioning pool capacity increases and is made available. Be aware that the pool might momentarily show a Normal (Formatting(0%)) status while it formats the newly created portion of the pool. Nevertheless, any previously data written to the pool is not affected and the new storage capacity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

## Virtual Server and Virtual Machine Configuration

Windows Hyper-V Manager allows for easy installation and configuration of your Hyper-V environment. This section provides general guidance with regards to configuring the overall Hyper-V environment as well as setting up your virtual machine for OS and SQL Server 2008 installation and configuration. However, this section is not intended to provide detailed setup and configuration of the OS and SQL Server 2008 installation and configuration because that falls outside of the scope of this implementation guide.

Ensure that all LUs to be utilized by the VM are readily available for use by the Hyper-V server. LUs need to be available for allocation to the VM's virtual hard disk and additional LUs for SQL Server database, transaction log, and tempdb files.

Also be sure to create your environment's virtual network before starting the virtual machine creation.

## Creating a Virtual Machine

Set up the location where the virtual machines' hard disks and configuration files will be stored. Hitachi recommends using the LU's GUID to identify volumes to be used as the virtual machine container. Using the GUID allows the Hyper-V host to bypass the 26 drive letter limitation in environments with a large number of storage volumes assigned to the server. Additionally, using GUID is extremely useful in cluster environment due to the manner in which volumes are identified within a cluster.

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 virtual machine, storage performance and capacity perspectives.

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

<i>Resource</i>	<i>Details</i>
Virtual machine	32GB memory 2 virtual processors
Hitachi Adaptable Modular Storage 2300 storage system	23 x 146GB 15K RPM SAS disks for SQL Server data files
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 1 x 62.5GB on RAID-1+0 (6D+6D) for tempdb 1 x 160GB on RAID-1+0 (2D+2D) for transaction logs 1 x 150GB on RAID-5 (4D+1P) for OS

OLTP workload testing conducted in the Hitachi Data Systems lab shows that two logical processors 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. Note that it's important to test your environment prior to deploying it in a production setting because your workload might differ from the one tested.

To configure the Adaptable Modular Storage 2300 LU to be used to store virtual disk and configuration file from the virtual machines and create your virtual machine using the newly created volume, follow these steps. Note that while these instructions describe how to identify a volume's GUID using Windows Disk Management, you can also acquire a volume's GUID with the `mountvol . exe` command.

1. Open Windows Server Manager and select **Disk Management** option under **Storage**.
2. Right click the disk to be used and select **Online**.
3. Right-click the same disk and select **Initialize**.
4. Right-click the disk's unallocated volume and select **New Simple Volume** to create a new partition on the volume within the disk.

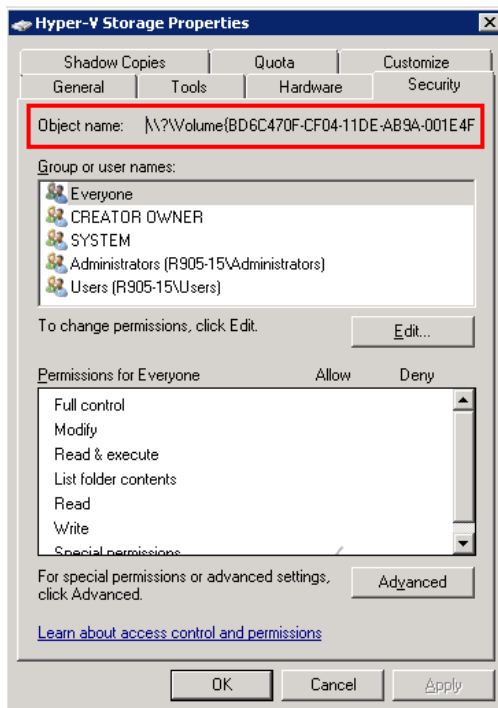
Use the New Simple Volume Wizard and make sure to select the **Do not assign a drive letter or path** option. However, give the volume a meaningful name such as “HyperV storage” to make it easily identifiable.

Note: When using Hitachi Dynamic Provisioning volumes, it is important to select the **Quick formatting** option. Not selecting the quick formatting option leads to the unnecessary use of storage space within a Dynamic Provisioning pools due to the nature of the full formatting execution.

The volume is created.

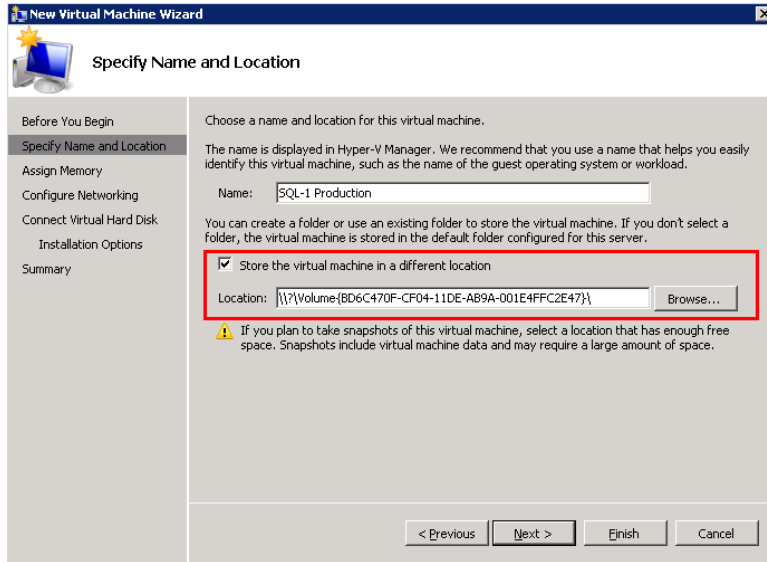
5. Right-click the volume and click **Properties**.
6. Select the **Security** tab and note the volume’s GUID (object name).
7. Select and copy that string.

You use the GUID to identify where the virtual machine files will be stored.



8. Open Hyper-V Manager and select the appropriate Hyper-V server.
9. On the right pane, click the **New** link and select **Virtual Machine** from the pop-up menu.
10. Click **Next** if the **Before You Begin** section is the first thing that appears and enter a name for the virtual machine being created.

11. Select the **Store the virtual machine in a different location** check box, enter the volume's GUID in the location field and click **Next**.

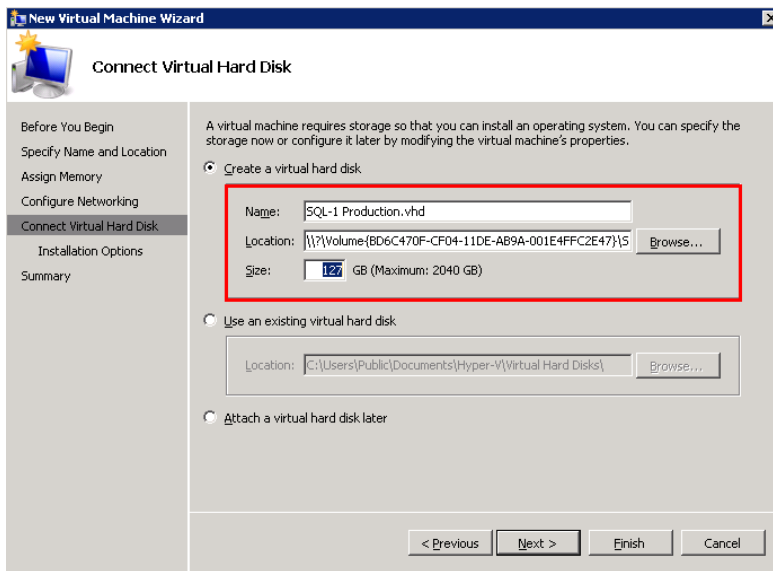


12. Assign memory to the virtual machine and click **Next**.

In this example, the virtual machine is assigned 32GB of memory (or 32,768MB).

13. Select the virtual network to which the virtual machine will be connected and click **Next**.

Notice that the **Location** field is automatically populated with the volume's GUID information and the virtual machine name entered in Step 10.



14. Enter a size for the virtual disk and click **Next**.

15. Select a VM operating system installation method and click **Next**.

16. Review all the selected information in the **Description** field, and if everything is correct, click **Finish**.

By default, only one virtual processor is assigned to a newly created virtual machine. To change this value, follow these steps:

- a. Select the desired virtual machine and click the **Settings** link in the right pane.
- b. Select the **Processor** option under the **Hardware** menu item.
- c. Select the appropriate number of logical processors from the drop-down menu and click **OK**.

Microsoft recommends starting with one logical processor and increasing the number as it becomes necessary.

When using a building block approach, Hitachi Data Systems recommends creating a single virtual machine and using it as a template to simplify deployment. You can also use SysPrep when preparing your virtual machine template, which can simplify the deployment of your virtual machines, although at this time, there is no support for SysPrep from a SQL Server 2008 perspective. Note that appropriate licensing for this task is still required from an OS and SQL Server perspective.

### *Assigning SQL Server Data Volumes to the Virtual Machine*

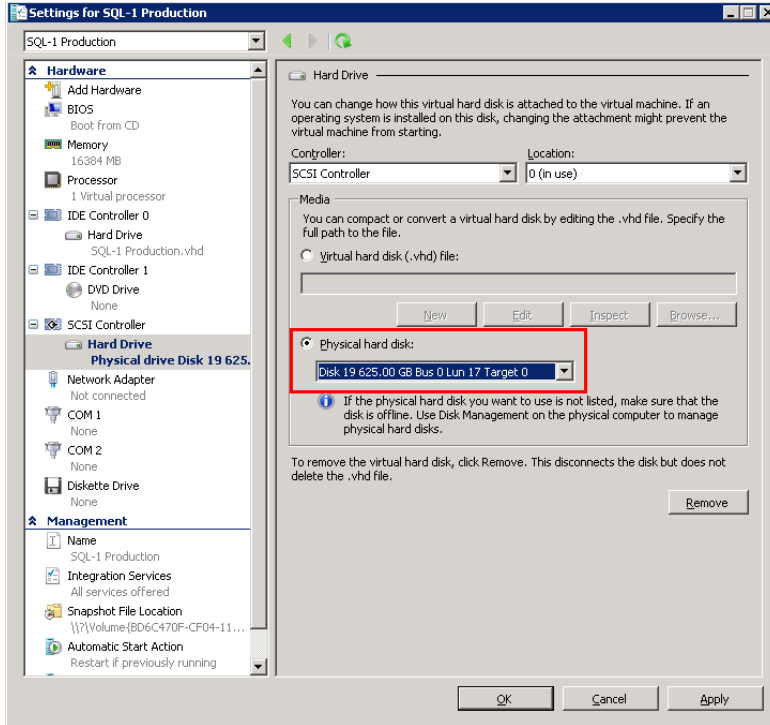
To assign volumes to a given virtual machine, you must first identify the volumes through the use of Server Manager on the Hyper-V server being used. Once identified, you must establish whether to use VHD files or SCSI pass-through disks when assigning the disks to the virtual machines. The solution described in this white paper uses pass-through disks. For more information about the options, see the [Virtualizing Microsoft® SQL Server on the Hitachi Adaptable Modular Storage 2000 Family Using Hyper-V Reference Architecture](#) white paper.

Follow these steps to identify and assign the data volumes to the virtual machine:

1. Open **Hyper-V Manager**.
2. Select a virtual machine and click the **Settings** link on the right pane.
3. Select the **SCSI Controller** option under the **Hardware** menu item.
4. Select the **Hard Drive** option and click **Add**.

5. Under the **Media** section, select the **Physical hard disk** option and select the disk that will be used for the SQL Server data files.

Notice that the disk number, size, and LUN information is displayed; ensure that the correct disk drives are selected.



6. Repeat Step 3 through Step 5 to add remaining disks to the virtual machine.
7. Click **OK**.

To bring the disk online, create a partition and format the volume for use with SQL Server, use the Server Manager from within the virtual machine. The disks can be either mounted using a drive letter or on a mount point. Because a SQL Server instance can scale up by housing multiple user databases requiring many LUs, Hitachi Data Systems recommends using mount points due to the drive letter limitation. For more information about standard Windows practices, see Windows online help.

For more information, see the [Hyper-V TechNet](#) Web site.

## 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 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 using tools to validate your SQL Server environment, see SQLCAT's "[Deploying SQL Server 2005 with SAN #3](#)" article.

You can download these tools from the following 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 and Microsoft 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 RAID groups and LUs to the environment or adding more Dynamic Provisioning LUs when the deployment uses Hitachi Dynamic Provisioning software. You can grow your environment in several ways:

- **Database size and I/O capacity** — Increase the capacity of your databases by adding additional RAID groups and LUs, or by simply creating additional Dynamic Provisioning LUs in your Dynamic Provisioning pool. Note that adding RAID groups to a given pool might also be necessary depending on how much of the pool capacity has already been used.
- **Virtual machines** — Increase the number of SQL Server hosts on your environment according to your hardware needs.
- **Hyper-V Servers** — Each Hyper-V server in this environment can accommodate a varying number of virtual machines. The limit is dependent on your individual virtual machine workload along with the assigned resources to each of the virtual machines.

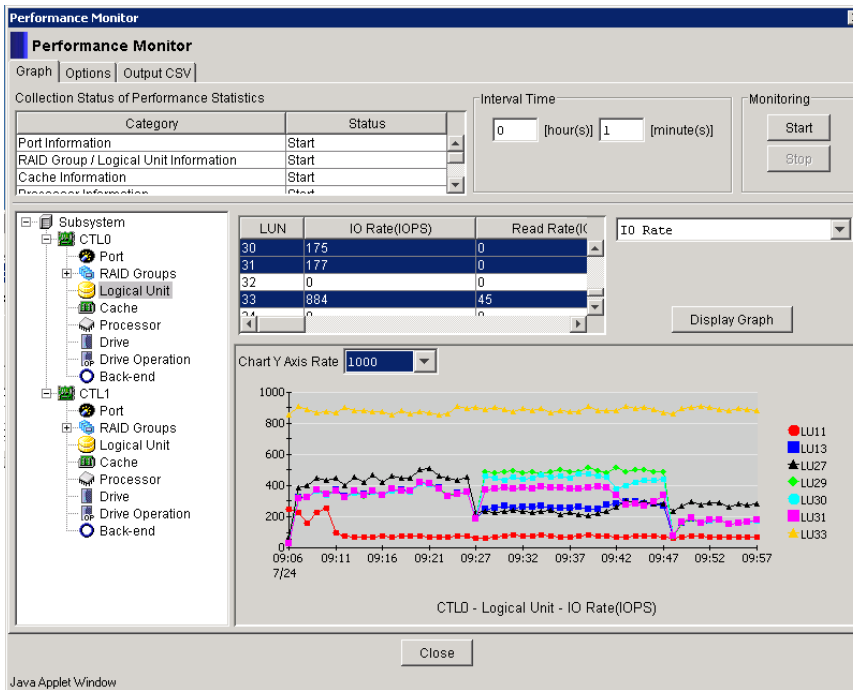
For this environment, the Hyper-V server contained 128GB of memory and four quad-core AMD Opteron CPUs. When planning to scale your environment, it's also important to consider your current size and the database workload expected from users to be added when deciding which increment is best. In all cases,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

### 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 4 shows the monitoring of an environment using Performance Monitor feature's GUI.

**Figure 4. Performance Monitor Feature User Interface**

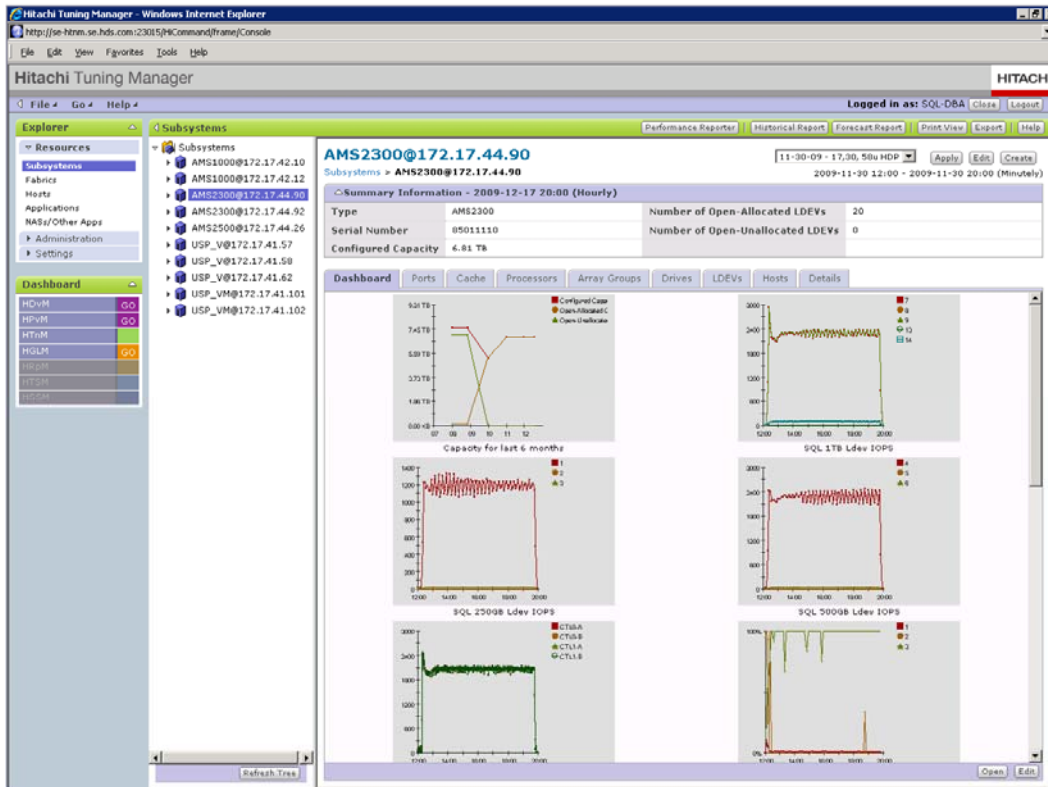


## Hitachi Tuning Manager Software

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 Hyper-V 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](#).

Figure 5 shows the monitoring of a SQL environment using Hitachi Tuning Manager software.


Figure 5. Hitachi Tuning Manager Software User Storage Dashboard



## Hyper-V Protection Strategies

A successful Hyper-V deployment requires careful consideration of protection strategies for backups, disaster recovery and quick migration, including the following:

- **Hitachi Storage Cluster Solution** - Integrating Hyper-V with Adaptable Modular Storage 2000 family replication solutions provides high availability for disaster recovery scenarios. This solution leverages the quick or live migration feature of Hyper-V to allow for the planned and unplanned recovery of child partitions under Hyper-V.
- **Hyper-V quick migration** - Hyper-V quick migration provides a solution for both planned and unplanned downtime. Planned downtime allows for the quick movement of virtualized workloads to service the underlying physical hardware. This is the most common scenario when considering the use of quick migration.
- **Hyper-V live migration** – Like Hyper-V quick migration, Hyper-V live migration provides a solution for both planned and unplanned downtime with the advantage of not having any disruption of service or perceived downtime. With Hyper-V live migration one can move running virtual machines between clustered Hyper-V servers without downtime.

- 
- **Backups** - Regularly scheduled backups of the Hyper-V servers and the data that resides on the child partitions under Hyper-V are an important part of any Hyper-V protection plan. With Hyper-V, the backup and protection process involves both the Hyper-V parent partition and the child partitions that execute under Hyper-V, along with the applications that reside within the child partition.
  - **Storage Replication** - Another important part of protection strategy is storage replication. The 2000 family has built-in storage replication features such as Hitachi ShadowImage Replication software and Hitachi Copy-on-Write software that can provide rapid recovery and backup in a Hyper-V environment.

For more information, see the [Microsoft Hyper-V Solutions page](#) on the Hitachi Data Systems Web site.



---

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

**Europe Headquarters** Sefton Park, Stoke Poges, Buckinghamshire SL2 4HD United Kingdom  
Contact Information: + 44 (0) 1753 618000 [www.hds.com](http://www.hds.com) / [info.uk@hds.com](mailto: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-034-01 March 2010