

Protecting Hyper-V CSVs with Microsoft Data Protection Manager and the Hitachi VSS Hardware Provider on the Hitachi Adaptable Modular Storage 2000 Family

Implementation Guide

By Rick Andersen

September 2010

Feedback

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

Table of Contents

- Solution Overview 2**
 - Understanding Protection for Cluster Shared Volumes 3
 - Protecting Virtual Machines using Hitachi VSS Hardware Provider 4
- Key Solution Components 6**
 - Hardware 6
 - Software 7
- Solution Configuration 9**
 - Storage Configuration 9
 - Microsoft Data Protection Manager 2010 Configuration 12
 - Hitachi VSS Hardware Provider Configuration 14
- Engineering Validation 17**

Protecting Hyper-V CSVs with Microsoft Data Protection Manager and the Hitachi VSS Hardware Provider and the Hitachi Adaptable Modular Storage 2000 Family

Configuration Guide

Backups are important, but it's just as important to perform backups without affecting user access to critical data during the backup window. Backup and recovery in a virtual environment are a complex endeavors, with several factors to consider. Backups can be done from the server running the virtualization software like Microsoft Hyper-V, or they can be done from within the guest virtual machines executing under the Hyper-V server. Using Microsoft Volume Shadow Copy Service (VSS) software and hardware providers are additional options that must be considered. Use of Cluster Shared Volumes (CSVs) in a Hyper-V environment where multiple VMs and their associated data files might be hosted on a single LU adds complexity.

The use of CSVs allows all nodes in a Hyper-V failover cluster concurrent access to data on a CSV-enabled disk. CSV implements I/O redirection techniques through a file system mini-filter driver. This redirection introduces new challenges to backup and restore operations that use the Hyper-V VSS writer for host-based backup and restoration of virtual machines. It is important to use an application that fully supports CSVs, such as Microsoft Data Protection Manager 2010, to ensure the integrity and usability of the backup. Support for CSV backups is provided by the Volume Snapshot Service (VSS); all CSV backup and restore operations must be done from within the VSS framework.

This white paper describes a solution that uses the Hitachi Adaptable Modular Storage 2000 family, Data Protection Manager 2010 and the Hitachi VSS Hardware Provider to help companies achieve those goals. The Hitachi VSS Hardware Provider is the key component of this solution; it integrates Hitachi storage hardware and software with VSS. VSS requires three components: a requestor, a provider and a writer. In this solution, the Hitachi VSS Hardware Provider functions as the provider, Data Protection Manager 2010 functions as the requestor and Hyper-V server functions as the writer.

This white paper provides guidance about dealing with some of the protection challenges that backup administrators face when protecting CSVs, especially the need to back up large quantities of data in short periods of time and to cost effectively meet strict recovery time objectives (RTOs) and recovery point objectives (RPOs). It describes how to configure Data Protection Manager 2010 to perform backups of CSVs in conjunction with using the Hitachi VSS Hardware Provider. This white paper also lists the hardware and software required to build the solution and provides links to supporting documentation needed to build, test and validate the solution. This paper also details how using the Hitachi VSS Hardware Provider vastly improves backup and restore times for virtual machines that are hosted on CSVs.

Although this document does not provide step-by-step, detailed instructions for every task and activity required to deploy the solution, it does serve as a resource where readers can easily locate related materials needed to build a functional solution. It is written for Windows and storage administrators charged with designing and implementing backup and recovery solutions. It assumes working knowledge of Data Protection Manager 2010 and Windows Server operating systems and a basic understanding of SAN-attached storage concepts.

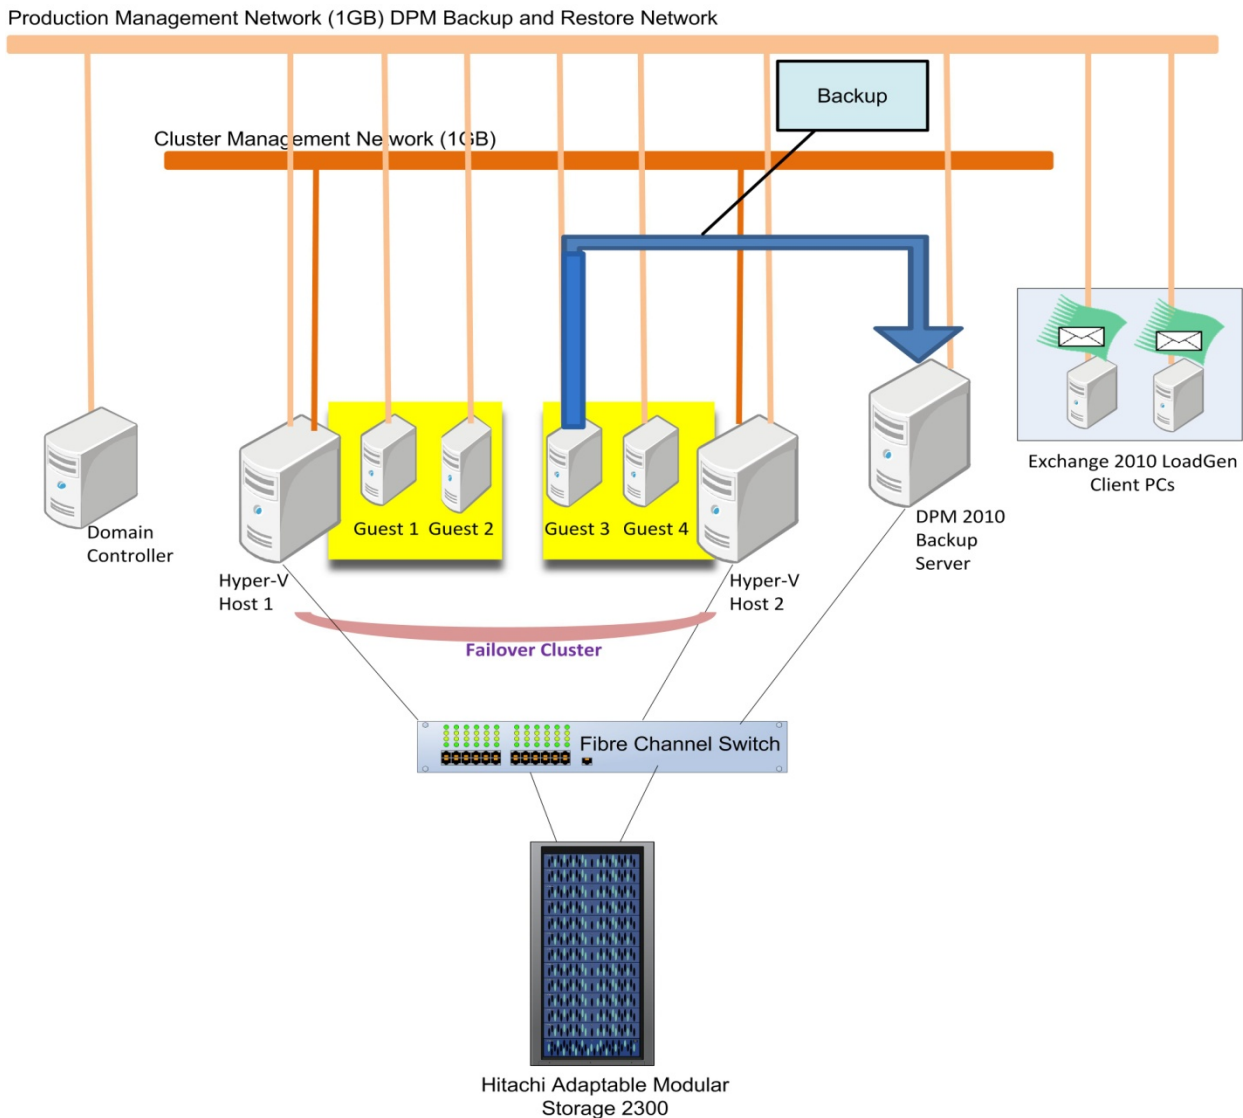
Hitachi Data Systems used a Hitachi Adaptable Modular Storage 2300 when testing this solution; however, any member of the 2000 family can be used to implement it.

Solution Overview

The solution described in this white paper assumes a functioning Hyper-V failover cluster environment and the use of CSVs to host the guest machines' OS and data files. This solution uses the two VSS providers available when backing up CSVs, the Hitachi VSS Hardware Provider and the VSS Software Provider that comes with Data Protection Manager 2010. Hitachi Data systems used Exchange Load Generator 2010 (LoadGen) on all four guest machine to generate a real world workload to test and validate the backup and restoration of CSVs.

Figure 1 illustrates the solution topology.

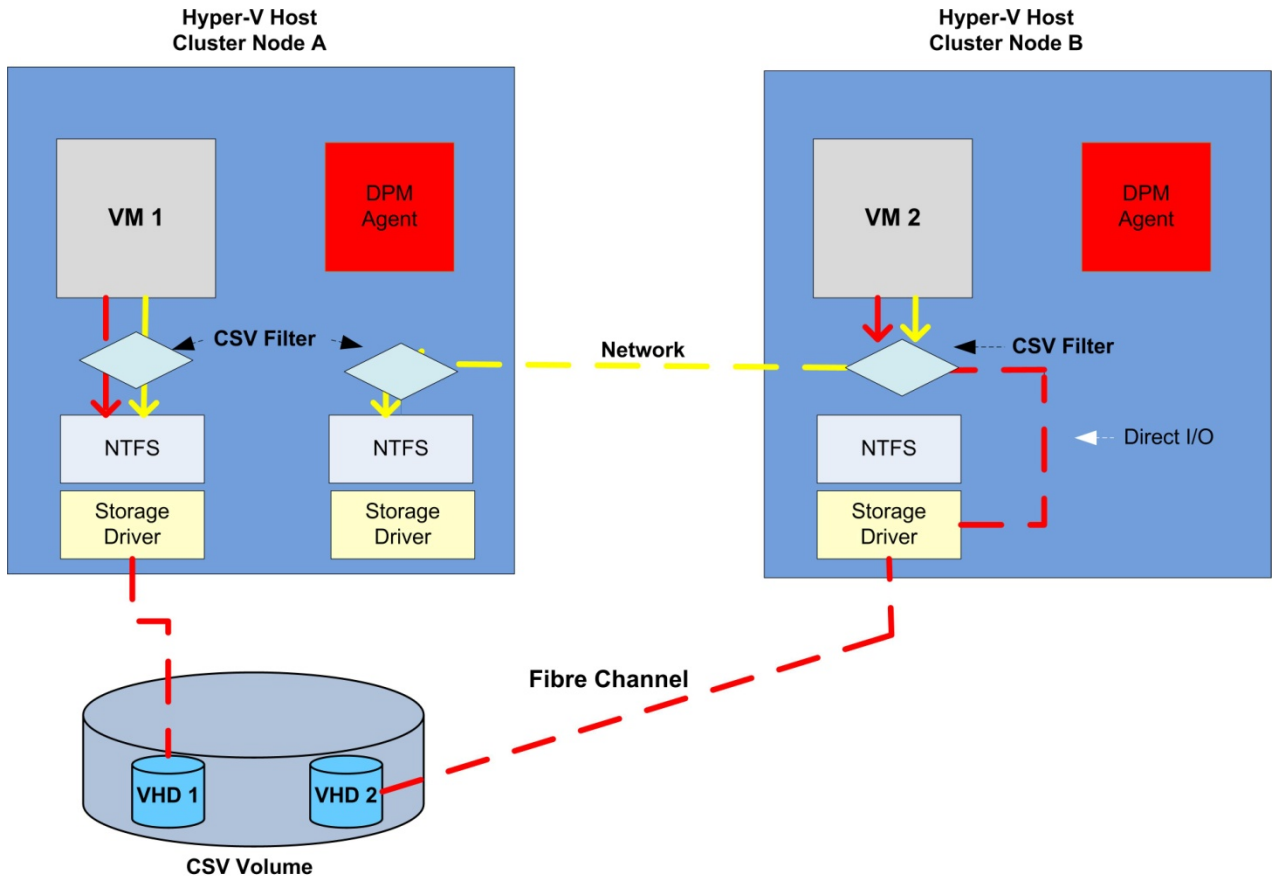
Figure 1. Solution Topology



Understanding Protection for Cluster Shared Volumes

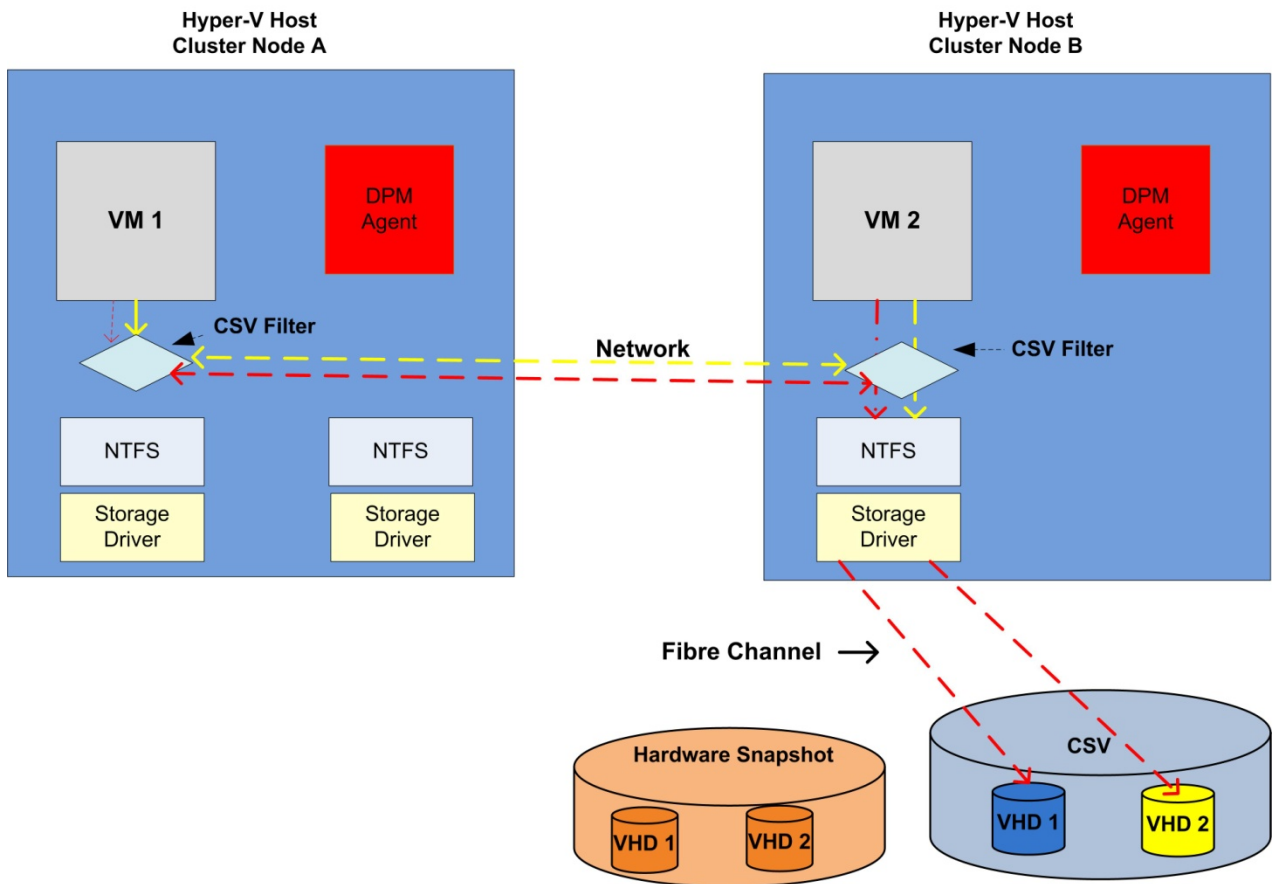
A standard CSV deployment places the VHDs for the virtual machines on a CSV with the virtual machines distributed across all the nodes in the cluster. Each virtual machine has direct I/O access to its respective VHD on the CSV. Figure 2 shows that, VM 2 has direct I/O access over a Fibre Channel network to its VHD on the CSV even though the CSV volume is owned by Cluster Node A.

Figure 2. Direct I/O Access



When you initiate a backup for VM 2 with Data Protection Manager 2010 (DPM), CSV ownership is moved to Cluster Node B. This causes all Fibre Channel I/O for VM1 to be routed over the network through the CSV filter on Node B. This affects the performance of VM1 because all I/O is now being redirected over the network, as shown in Figure 3.

Figure 3. Redirected I/O Access



To reduce the negative performance effect of redirected I/O on VM 1, use the Hitachi VSS Hardware Provider, which allows the CSV to resume direct I/O as soon as the hardware snapshot is taken. Use of the Hitachi VSS Hardware provider shortens the duration of the redirected I/O to approximately two minutes in most cases.

If you use the software provider instead of the hardware provider to create backups, the CSV is in redirected I/O mode for all the virtual machines on the CSV for the duration of the backup. This backup duration depends on the size of the VHD being backed up by DPM 2010 and the time can be significant. During the duration of redirected I/O mode, the performance of other virtual machines on the CSV is reduced.

For more information about the use of the Hitachi VSS Hardware Provider and the software provider, see the "Engineering Validation" section of this paper.

Protecting Virtual Machines using Hitachi VSS Hardware Provider

This solution includes three phases:

1. Pair and split
2. Backup
3. Restore


The Hitachi ShadowImage® In-System Replication software pair and split phase, facilitated by Hitachi VSS Hardware Provider, reduces the backup window to only a matter of seconds per virtual machine being backed up. The Hitachi VSS Hardware Provider works with the Microsoft Volume Shadow Copy service to create a clone, or a full copy, of the original data on a volume. The data on this clone can be backed up without disrupting the production system.

To backup virtual machines and their respective files on a CSV using VSS, the following sets of LUs are required:

- Primary volumes (P-VOLs) for each guest machines system files, databases and transaction logs
- Secondary volumes (S-VOLs) for each guest machines system files, databases and transaction logs
- LUs to create the DPM backup pools that will house the backed up data or, as alternative, tape libraries or virtual tape libraries
- Two small LUs for the differential management logical units (DMLUs) used by ShadowImage software

Figure 4 shows the drive layout for the tested implementation of the backup and recovery solution described in this configuration guide.

Figure 4. Drive Layout



For more information about the tested deployment, see the “Storage Configuration” section of this paper.

Key Solution Components

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

Hardware

For this solution, Hitachi Data Systems connected the Hyper-V host servers, Data Protection Manager 2010 and the Hitachi 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.

Table 1 describes the storage and SAN components used in the tested implementation.

Table 1. Tested Deployment Hardware

<i>Hardware</i>	<i>Configuration</i>
Hitachi Adaptable Modular Storage 2300 storage system	Microcode 0880 A/M* 2 storage controllers 16GB cache per controller 7 disk trays (RKA), 4 trays used 15 x 300GB 15K RPM SAS disks per tray
Brocade 300 SAN switch	FOS 5.3.1a* 3 4Gb Fibre Channel ports used
Hitachi BladeSymphony 2000	2 Intel Xeon E5520 processors 32GB memory 4 Hitachi HFC0402-E Fibre Channel HBAs 4Gb/s Hyper-V server
Hitachi BladeSymphony 320	2 Intel Xeon E5570 processors 32GB memory 1 Hitachi HFC0402-E Fibre Channel HBAs 4Gb/s Microsoft Data Protection Manager 2010 server

*Use the latest Hitachi supported levels of microcode and FOS levels.

Hitachi Adaptable Modular Storage 2000 Family

The storage systems that make up the Hitachi Adaptable Modular Storage 2000 family are the first midrange products to offer a serial attached SCSI (SAS) architecture and the Hitachi Dynamic Load Balancing Controller. The 2000 family delivers highly resilient, enterprise-quality storage in an affordable and easy-to-manage modular package.

The Hitachi Adaptable Modular Storage 2300, which was used for testing this solution, offers best combination of price and performance in a model that scales to 240 disk drives. Ideal for large businesses and enterprises, the 2300 is a highly reliable, flexible and scalable storage system for mission-critical business applications.

Although Hitachi Data Systems used the 2300 in its lab, any member of the 2000 family is appropriate for this solution.

Hitachi BladeSymphony

Hitachi BladeSymphony servers are powerful, flexible and reliable. The BladeSymphony 2000 is a 10U, eight-blade system that handles the most I/O-intensive workloads with ease. It features a balanced system architecture that combines massive throughput, embedded virtualization for consolidating systems and workloads, and unprecedented configuration flexibility. The BladeSymphony 320 is the industry's first blade server with a 110-volt power option, and it is lightweight and easy to move and reconfigure. Yet this 6U system packs up to 840 cores into a standard 42U rack, a space savings of 60 percent compared to rack servers.

Software

The following sections describe the software used to deploy this solution.

Table 2 lists all of the software used in testing this solution.

Table 2. Tested Deployment Software

<i>Software</i>	<i>Version</i>
Windows Server (for Hyper-V server)	2008, R2, Enterprise
Windows Server (for all other servers)	2008, R2, Enterprise
Exchange Server	2010, Enterprise
Hitachi Storage Navigator Modular 2	7.0
Microsoft Data Protection Manger 2010	RTM
Hitachi VSS Hardware Provider	4.0.1

Hitachi Dynamic Link Manager Software

This solution used Hitachi Dynamic Link Manager software's round-robin multipathing policy and configured two redundant paths from the servers to the 2300. Each server had two host bus adapters (HBAs) for high availability purposes. Hitachi Data Systems 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. Another option is to use the round-robin algorithm available with Windows native MPIO feature.

Hitachi ShadowImage In-System Replication Software

This solution used Hitachi ShadowImage In-System Replication software, which uses local mirroring technology, to create full-volume copies or clones within the Hitachi Adaptable Modular Storage 2300. Although ShadowImage software is the underlying technology that replicates the volumes necessary to perform a backup from the secondary volume, it cannot be used to achieve consistent, point-in-time backups of virtual machines on CSVs without integrating into the Microsoft VSS framework. This integration occurs through the Hitachi VSS Hardware Provider. For more information, see the "Hitachi VSS Hardware Provider" section in this white paper.

Hitachi VSS Hardware Provider

The Hitachi VSS Hardware Provider works with Microsoft's Volume Shadow Copy Service (VSS) to generate consistent point-in-time copies of data known as shadow copies. Microsoft's VSS was introduced in Windows Server 2003, and Windows Server 2008 contains an enhanced version of the VSS framework. VSS works with three main components to produce the shadow copies. Table 3 lists definitions of these components.

Table 3. Required Volume Shadow Copy Service Components

Component	Description
Requestor	Application, such as Microsoft Data Protection Manager 2010, that requests that a volume shadow copy be taken
Writer	Software that is included with applications or an OS, such as Hyper-V, that helps provide consistent shadow copies
Provider	Component that creates and maintains the shadow copies

For Hyper-V host backups of guest machines to occur on the Hitachi Adaptable Modular Storage 2300 using its internal storage replication capabilities, the Hitachi VSS Hardware Provider must be installed on both the Hyper-V host and Data Protection Manager backup server. Hardware-based shadow copy providers, like Hitachi VSS Hardware Provider, act as an interface between Microsoft VSS and the storage system. The work of creating a shadow copy is done by the 2000 family storage controller; Microsoft VSS ensures that the copy and other processes comply with VSS standards. Download the Hitachi VSS Hardware Provider from the Hitachi Data Systems [web site](#).

Microsoft Data Protection Manager 2010

Data Protection Manager 2010 provides continuous protection data protection for virtual machines hosted on servers running Microsoft Hyper-V. This protection includes online backup of guest virtual machines hosted on a clustered or standalone environment, protection of virtual machines during the live migration process and item-level recovery from a host-level backup. DPM 2010 interacts with the Hyper-V VSS writer so that consistent versions of a virtual machine can be captured and protected without impacting client access to the virtual machine.

Table 4 lists the DPM 2010 components needed to protect virtual machines under Hyper-V.

Table 4. DPM Components

Component	Description
DPM 2010 server	DPM 2010 server performs the replication synchronization, and recovery point creation for protection and recovery of data.
DPM 2010 agent	Protection agents must be installed on each Hyper-V server hosting virtual machines requiring protection. This agent tracks changes to protected data and transfers the changes to the DPM 2010 server.
DPM 2010 protection groups	Protection groups are used to manage the protection of data on virtual machines.
SQL Server 2008 SP1	DPM 2010 requires a dedicated instance of SQL Server 2008, 32-bit or 64-bit version, and Enterprise or Standard edition.

Once installed and configured, these components can be controlled and monitored from the DPM 2010 console.

Solution Configuration

To implement this solution on a 2000 family storage system, follow these high-level steps:

1. Configure the 2000 family storage system.
2. Install and configure the Microsoft Data Protection Manger 2010 server.
3. Install the DPM 2010 agents on those servers requiring protection.
4. Install the Hitachi VSS Hardware Provider on the Hyper-V Host servers and the Microsoft DPM 2010 server.
5. Configure DPM 2010 to use the Hitachi VSS Hardware Provider.
6. Create the ShadowImage secondary volumes.
7. Create protection groups on the DPM 2010 server.

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

The specific activities and detailed processes for each of these high-level tasks are located in documentation provided by Hitachi and Microsoft. This documentation set is required to assist with deploying the solution. For more information, see the following resources:

- Hitachi Storage Navigator Modular 2 online help
- Hitachi VSS Hardware Provider software companion guide
- Hitachi ShadowImage In-System Replication User Guide that accompanies the software
- Microsoft Data Protection Manager 2010 online help

Storage Configuration

The configuration used for the tested implementation consisted of a two-node Hyper-V failover cluster hosting four virtual machines each executing an Exchange workload using the Exchange 2010 Load Generator. Each Exchange Server instance hosted 1000 mailboxes. For more information about the workload parameters used for the Exchange Workload Generator, see the “Engineering Validation” section of this white paper. This architecture also implements Hitachi Dynamic Provisioning pools, which take advantage of Hitachi Dynamic Provisioning software’s performance-enhancing, wide-striping feature for both the Exchange database and transaction log volumes and the secondary ShadowImage copies for those volumes.

Table 5 describes the tested storage configuration.

Table 5. Tested Storage Configuration

<i>Pool</i>	<i>RAID Level</i>	<i>Number of RAID Groups</i>	<i>Number of LUs</i>	<i>LU Size (GB)</i>	<i>Role</i>
1	1+0 (2D+2D)	3	1	800	CSV #1 contains the system VHD and Exchange database VHD for each guest
2	5 (4D+1P)	1	1	560	CSV # 2 contains the Exchange log VHD for each guest
3	1+0 (2D+2D)	3	1	800	ShadowImage S-VOL for CSV #1
4	5 (4D+1P)	1	1	560	ShadowImage S-VOL for CSV #2
5	5 (4D+1P)	2	1	2000	DPM backup pool

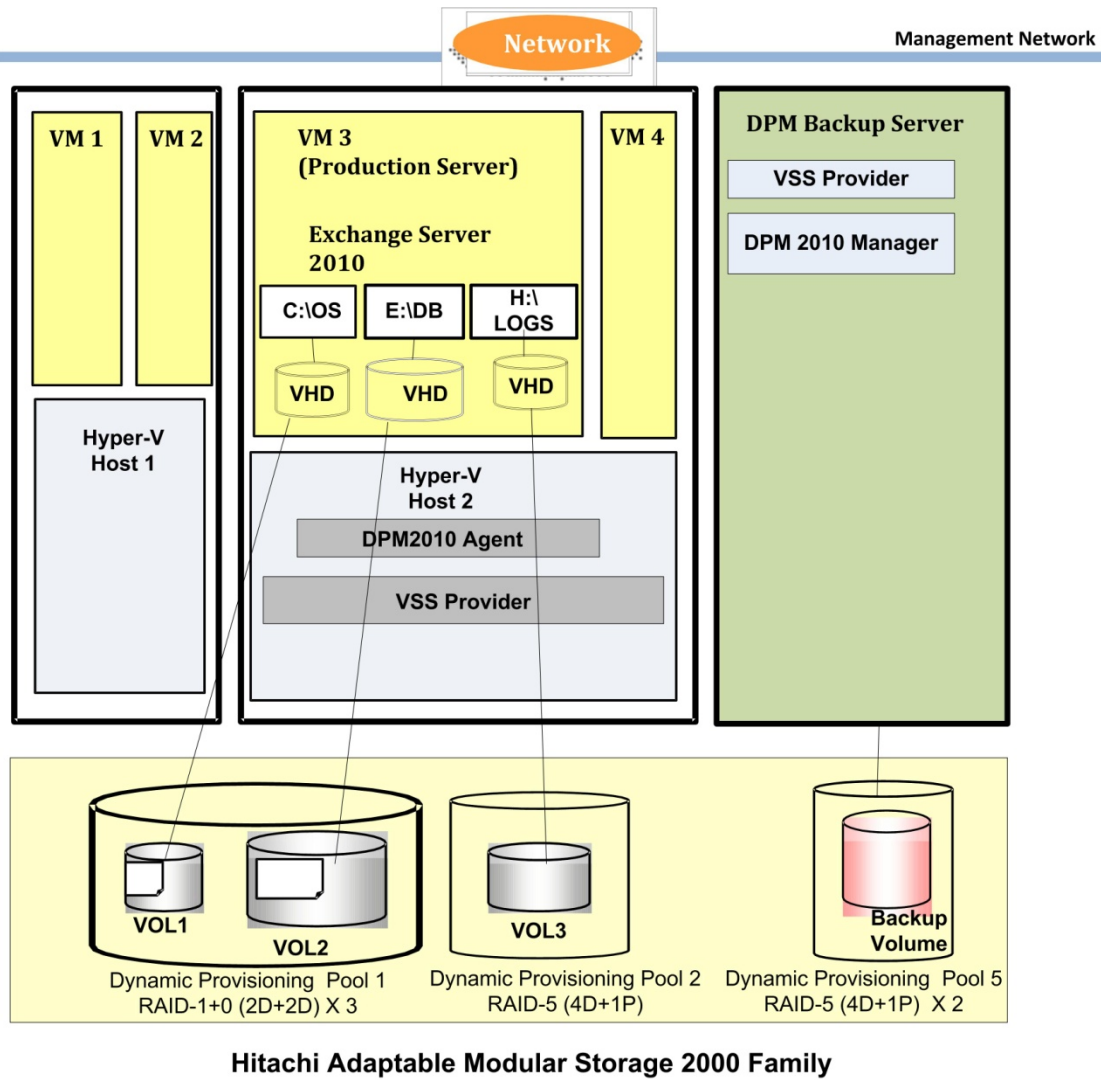
Table 6 describes how the VHDs for each guest machine were allocated on the CSVs.

Table 6. VHD to CSV mapping

<i>CSV LU</i>	<i>CSV Path</i>	<i>VHD</i>	<i>Size (GB)</i>	<i>Mapping</i>
1	C:\ClusterStorage\Volume1	System 1	50	Guest #1
		Mailbox Database 1	150	
		System 2	50	Guest #2
		Mailbox Database 2	150	
		System 3	50	Guest #3
		Mailbox Database 3	150	
		System 4	50	Guest #4
		Mailbox Database 4	150	
2	C:\ClusterStorage\Volume2	Transaction Log 1	75	Guest #1
		Transaction Log 2	75	Guest #2
		Transaction Log 3	75	Guest #3
		Transaction Log 4	75	Guest #4

Figure 5 shows the mapping of the storage for this solution from the Hyper-V host to the 2300 using the example of the mapping of VM 3.

Figure 5. Storage Mapping



Use Hitachi Storage Navigator Modular 2 software to configure the Dynamic Provisioning pools, RAID groups and LUs. You create secondary volumes for the Exchange storage groups and logs later using the Hitachi VSS Hardware Provider GUI.

Differential Management Logical Volumes

A differential management logical unit (DMLU) is an exclusive volume used for storing ShadowImage information. Configure DMLUs before using ShadowImage software. The DMLU is like other volumes in the 2000 family storage system, but is hidden from a host. The minimum size for a DMLU is 10GB. Although only one DMLU is needed, Hitachi Data Systems recommends using two for redundancy. Hitachi Data Systems recommends creating the DMLUs on separate RAID-5 groups.

Creating a DMLU

To create a DMLU, follow these steps:

1. In Storage Navigator Modular 2 software, create a standard 10GB LU.

This 10GB LU is used to create a DMLU.

2. In the **Arrays** navigation tree, choose **Settings > DMLU**.

The **Differential Management Logical Units** pane displays.

3. Click the **Add DMLU** button.

The Add DMLU window displays.

4. Select the check boxes for the LUs that you want to assign as DMLUs and click **OK**.

A confirmation message displays.

5. Select the **Yes, I have read** check box and click the **Confirm** button.

A success message displays.

6. Click the **Close** button.

Microsoft Data Protection Manager 2010 Configuration

The Data Protection Manager 2010 server must first be installed along with SQL Server 2008 SP1. DPM setup installs a dedicated SQL Server instance on the DPM server or you can specify a remote SQL Server instance.

Define a storage pool for the DPM server. The storage pool is a set of disks that store replicas and recovery points for protected data. The capacity of the storage pool that you assign to the DPM server must be sufficient to provide disk-based protection of the selected data sources. This solution uses LUs from the 2300 for the DPM server storage pool. The DPM 2010 Hyper-V storage calculator is available from Microsoft to help plan the storage pool requirements for DPM. This calculator provides guidance on three aspects of data protection for Hyper-V machines:

- Estimates how much storage is required to protect the Hyper-V environment
- Estimates backup windows based on times for initial replication and incremental synchronization
- Outlines the type and class of server, memory requirement, and other system requirements required to create the DPM server

For more information about defining storage pools for the DPM server, see the Microsoft TechNet article, ["Planning a DPM 2010 Deployment."](#)

You must install the DPM Protection Agent on the Hyper-V hosts that are to be protected. You can use the Protection Agent Installation in DPM server to install the agent. Pass-through and iSCSI disks cannot be protected through the DPM agent installed in the Hyper-V host. In this solution, guest machines are hosted on VHDs.

Configuring DPM 2010 Protection Groups

Protection Groups in DPM 2010 allow you to group similar resources so that they can be protected in a similar way. Virtual machines in the same protection group use the same protection method and use the same short-term and long-term protection policies.

To create a protection group, follow these steps:

1. Launch Microsoft Data Protection Manager 2010 software.
2. In the DPM 2010 **Administrator Console Action** pane, click the **Create Protection Group** link.

The **New Protection Group** wizard launches.

3. Click the **Next** button.

The **Select Protection Group Type** dialog box displays.

4. Select the **Servers** radio button and click the **Next** button.

The **Select Group Members** dialog box displays.

5. Expand the **Available members** directory in the navigation tree, choose the data to protect by selecting the check boxes for those virtual machines you want to protect and click **Next**.

The **Select Data Protection Method** dialog box displays.

6. Enter a name in the **Protection group name** field, select the **I want short-term protection using Disk** check box and click **Next**.

The **Specify Short-Term Goals** dialog box displays.

7. Determine a retention range in days and select the range in the **Retention range** box.
8. Click the **Modify** button.

The **Express Full Backup Dialog** box displays.

9. From the drop-down menu, choose the day and time that the backup will run and click **Next**.

The **Review Disk Allocation** dialog box displays and shows the amount of disk space allocated for the protection group.

10. Select the **Automatically grow the volumes** check box and click **Next**.

The **Choose Replication Method** dialog box displays.

11. Select the **Automatically over the network** radio button, select either the **Now** radio button or the **Later** radio button and click **Next**.

The **Consistency check options** dialog box displays.

12. Select the **Run a consistency check for a replica if it becomes inconsistent** radio button and click the **Next** button.

The **Summary** dialog box displays.

13. Click the **Create Group button**.

The **Status** box displays showing that the protection group was successfully created.

As part of the process of creating a protection group, a full replica backup of the virtual machines specified as part of the protection group is taken. From this point on, *only* incremental backups occur based on the parameters specified in Step 7.

For more information about creating protection groups, see the Microsoft TechNet article "[Creating Protection Groups.](#)"

Hitachi VSS Hardware Provider Configuration

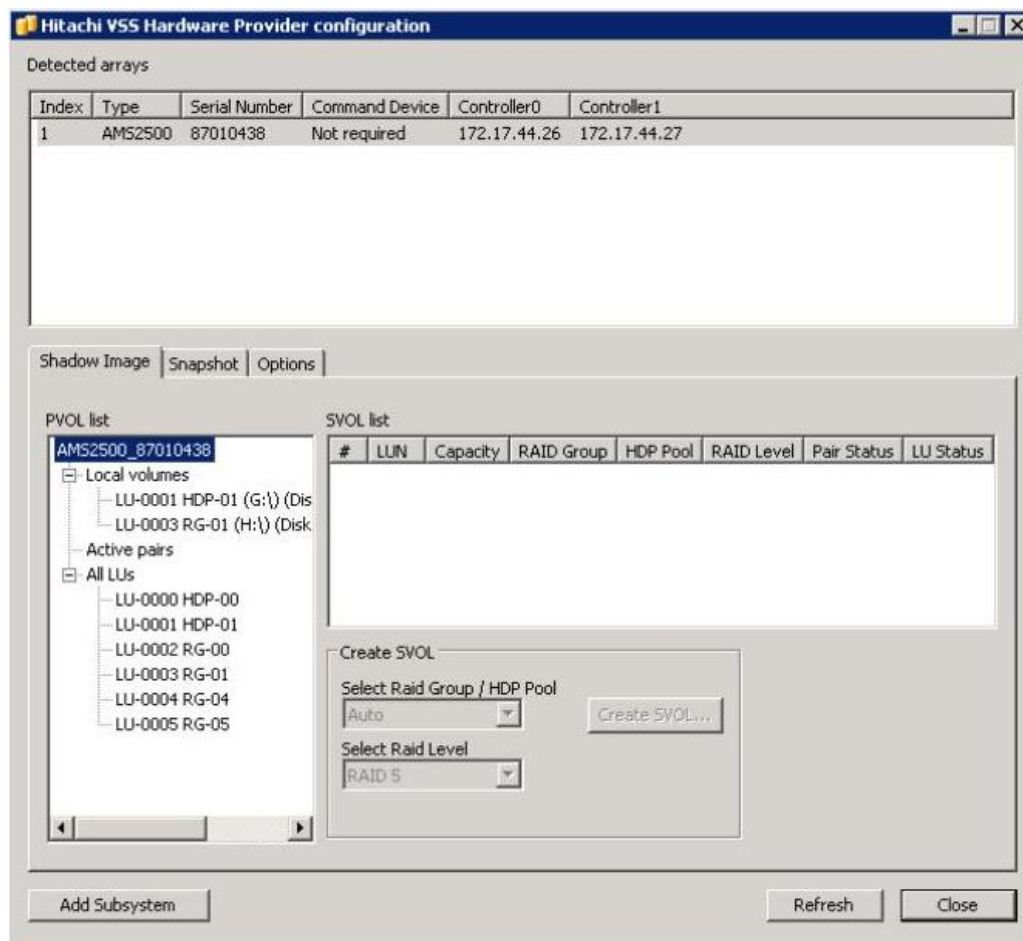
Before a backup can be taken, the ShadowImage secondary volumes or S-VOLs must be created and paired with the primary volumes or P-VOLs. Use the Hitachi VSS Hardware Provider to create the S-VOLs.

Adding a Storage System

To add a storage system to the Hitachi VSS Hardware Provider GUI, follow these steps:

1. In the Windows **Start** menu, click **Programs > Hitachi > VSS Hardware Provider > VSS Hardware Provider configuration.**

The VSS Hardware Provider GUI launches.



2. Click the **Add Subsystem** button.

The **Add Subsystem Login** dialog box displays.

3. Enter the IP address for controller 0 in the **Subsystem CTRL IP #0** field, the IP address for controller 1 in the **Subsystem CTRL IP #1** field and click **OK**.

The 2000 family storage system is displayed at the top of the window.

Configuring DPM 2010 to use the VSS Hardware Provider

To configure the DPM 2010 server to use the VSS Hardware provider, follow these steps:

1. Install the Hitachi VSS Hardware provider on the Hyper-V Host computers and the DPM server.

For more information, see the “Hitachi VSS Hardware Provider Configuration” section of this paper.

2. Delete the %ProgramFiles%\Microsoft DPM\DPM\Config\DataSourceGroups.xml file from the host computer.

Creating ShadowImage Pairs

Creating ShadowImage pairs establishes the replication relationship between the LUs that make up the P-VOLs and S-VOLs. When this relationship is established, ShadowImage places the pairs into a simplex (SMPL) state. The initial replication process involves a bit-for-bit copy from P-VOL to S-VOL, at which time the status of the pair is COPY. When the initial copy process is complete, the pair transitions into a PAIR state and is ready for backup through the VSS process.

To create ShadowImage pairs, follow these steps:

1. Launch the VSS Hardware Provider GUI.
2. Highlight the P-VOL under the **Local Volumes** category for which you want to establish an S-VOL.
3. In the **Create SVOL** pane, choose a RAID group or Dynamic Provisioning pool from the **Select Raid Group/ HDP Pool** drop-down menu.

Choose a RAID group or Dynamic Provisioning pool that was preconfigured to be used for S-VOL creation. The **RAID Level** field is automatically populated when the RAID group or Dynamic Provisioning pool is selected.

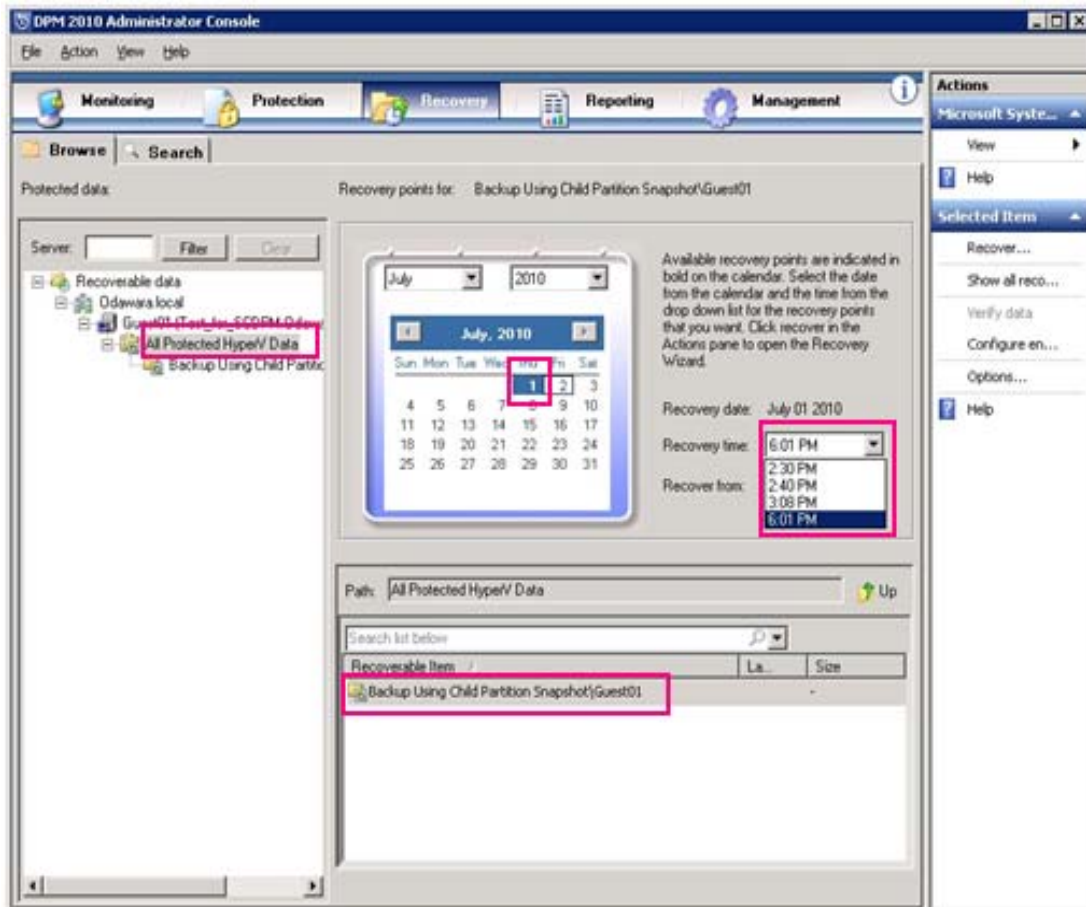
4. Click the **Create SVOL** button.

An S-VOL LU on the 2000 family storage system is created and puts the two LUs into a PAIR state.

Restoring a Virtual Machine in a CSV environment

To restore a virtual machine, follow these steps:

1. In the DPM Administrator console, click **Recovery** on the navigation bar.
The **Recovery points** dialog box displays.
2. Browse for the virtual machine listed under the cluster node that you want to recover.



Available recovery points are indicated in bold on the calendar in the recovery points section. Select the bold date for the recovery point you want to recover.

3. In the **Actions** pane, click the **Recover** button.
DPM's Recovery Wizard launches.
4. Review your recovery selections and click **Next**.
The **Select Recovery Type** dialog box displays.

5. Select the type of recovery you would like to perform.

You can choose to recover the original instance, recover to network folder, or recover to an alternate location.

6. Specify your recovery options and click **Next**.

The **Specify Recovery Options** box displays.

Note: Do not select the **Enable SAN based recovery using hardware snapshots** check box.

7. Click **Next**.

The **Summary** dialog box displays.

8. Click the **Recover** button to start the recovery operation.

Engineering Validation

To validate the functionality and the performance of backing up and restoring CSVs with Microsoft Data Protection Manager 2010 using the Hitachi VSS Hardware Provider, Hitachi Data Systems executed the following functional tests:

- Host-level VSS software backup
- Host-level VSS hardware backup using Hitachi VSS Hardware Provider with ShadowImage

For both test scenarios Exchange 2010 LoadGen was used on all four guest machine to generate a real world workload, which allowed Hitachi Data Systems to test and validate the backup and restore of CSVs in a real-world scenario.

Table 7 lists the backup attempts for each scenario tested and the results.

Table 7. Tested Scenarios and Results

<i>Scenario Description</i>	<i>Result</i>	<i>Comments</i>
Host-level VSS Software backup	Success	Backup much slower than VSS Hardware Provider.
Host-level VSS Hardware Provider backup (ShadowImage)	Success	The backup is much faster due to the replica being created from the S-VOL, which avoids contention with the production volume's I/O. Also reduces negative effect on other guest machines on the CSV volume to shortened duration of I/O redirection.

Hitachi Data Systems tested the Hitachi VSS Hardware Provider to determine the functionality and performance characteristics of protecting CSVs on the 2000 family, and tested the VSS Software Provider to protect CSVs on the 2000 family.

Table 8 shows the performance comparison between using the Hitachi VSS Hardware Provider and the VSS Software Provider to create the initial replica. All tests used 320GB VHDs.

Table 8. Initial Replica

<i>Item</i>	<i>MB/s</i>	<i>Time (Minutes)</i>
VSS Hardware Provider backup (ShadowImage)	50.8	105
Host-level VSS software backup	19.8	275

Table 9 shows the results of performing an express backup. An initial replica is the initial backup of a VHD. An express backup is a type of synchronization in which the DPM agent transfers a snapshot of all blocks that have changed since the initial replica backup or the last express backup. The impact of an express backup operation is expected to be less than the impact of a full backup since only changed blocks are transferred.

Table 9. Express Backup

<i>Item</i>	<i>Changed Blocks (GB)</i>	<i>MB/s</i>	<i>Time (Minutes)</i>
VSS Hardware Provider backup (ShadowImage)	7	2.3	50
Host-level VSS software backup	14	1.3	180

Hitachi Data Systems also measured the network utilization when performing a VSS Host level software backup in comparison to doing a VSS hardware backup. Figure 6 shows that network utilization rises and stays at a high level for the length of the VSS software backup, negatively affecting network performance.

Figure 6. Backup Using VSS Software Provider

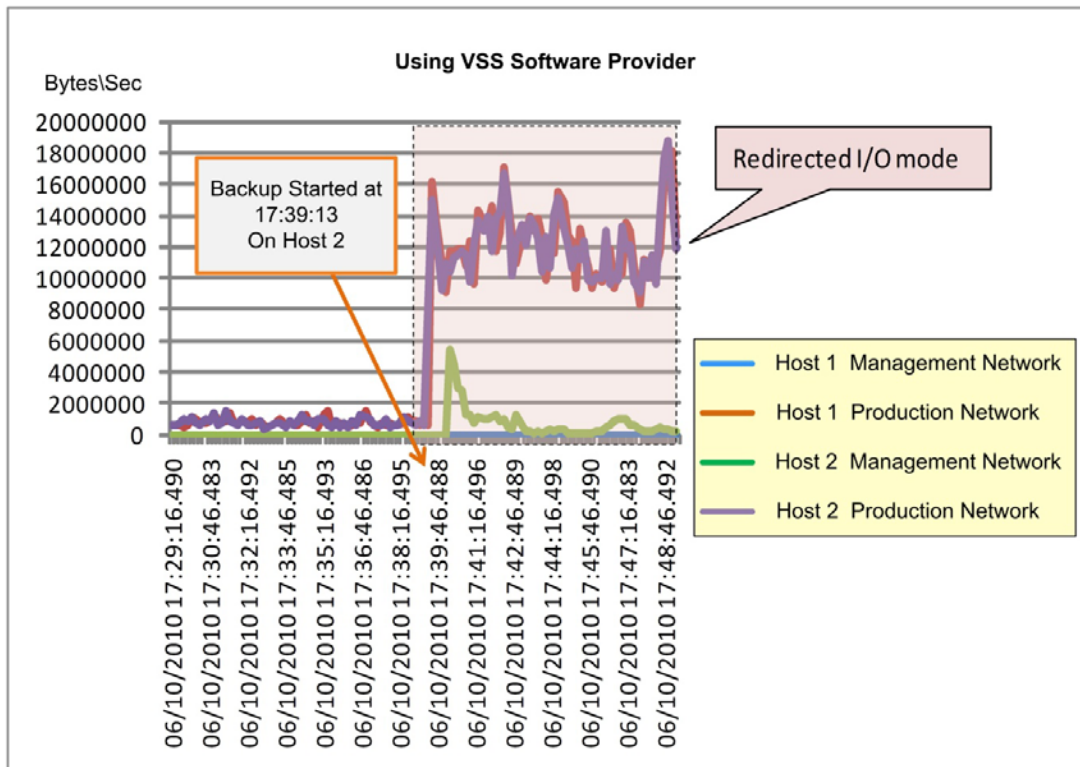
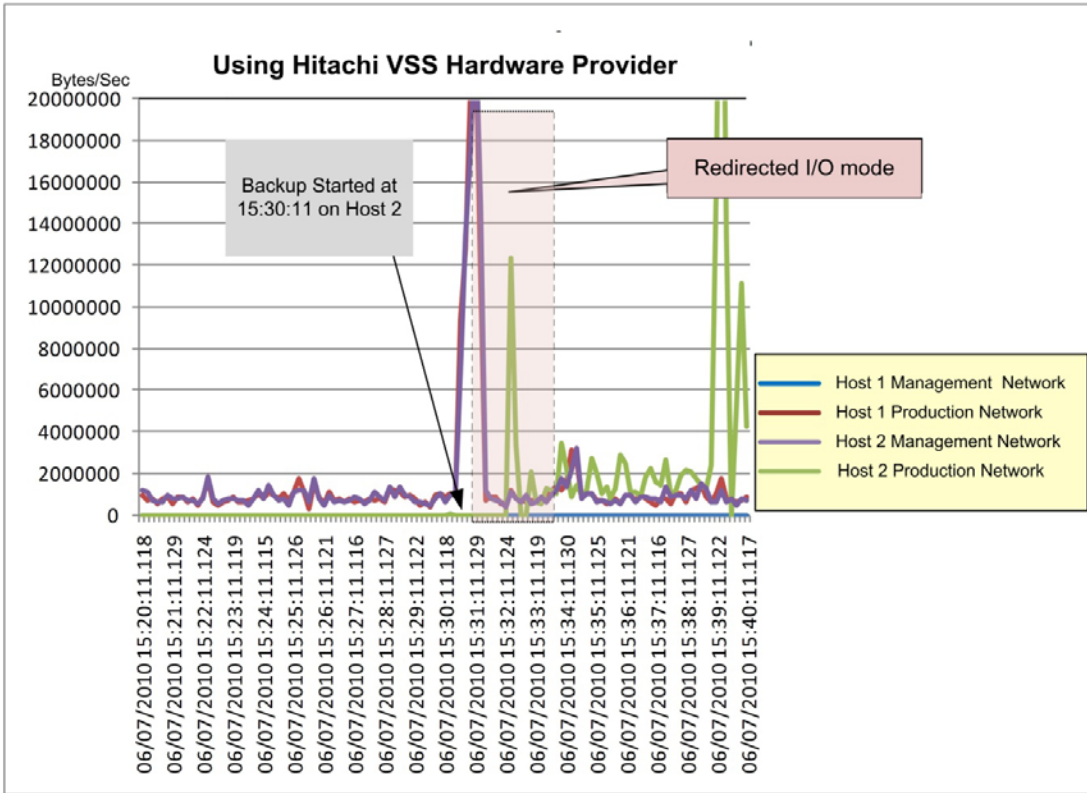


Figure 7 shows how the Hitachi VSS Hardware Provider greatly shortens the time that I/O to the CSV is in redirected mode over the network and thus reduces the network load.

Figure 7. Backup Using the Hitachi VSS Hardware Provider



Exchange LoadGen Parameters

Table 10 lists the parameters used as input to Exchange 2010 LoadGen to simulate an Exchange workload on all four guests during the backup and restore scenarios.

Table 10. LoadGen Parameters

<i>Parameter</i>	<i>Value</i>
Number of users per database	1,000
Average size of Mailbox per user	100MB
Average number of mail transfers per user	100 (send 20, receive 80)
Average size of mail transfer	100KB
Outlook Profile	Outlook 2007 cached
User Profile	Heavy



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 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-057-00 September 2010