Technical Report

NetApp SnapManager 2.0 for Hyper-V on Data ONTAP Operating in 7-Mode: Best Practices Guide

Santhosh Harihara Rao, NetApp
October 2013 | TR-4234

Abstract

This technical report provides guidelines and best practices for integrated architectures and implementations of Microsoft® Hyper-V® with NetApp® storage solutions. The NetApp technologies discussed in this technical report are important to achieving an integrated storage solution that is cost effective, operationally efficient, flexible, and environment friendly.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 SnapManager 2.0 for Hyper-V

With the adoption of virtualization technologies, data centers have been transformed and the number of physical servers drastically reduced. Virtualization has had many positive effects, not only reducing the number of physical systems, but also reducing network, power, and administrative overhead.

In contrast to physical environments, where server resources are underutilized, fewer resources are available in virtualized environments. Although each physical server had dedicated network and CPU resources, VMs must now share those same resources, which can result in performance issues, especially while backing up the virtual environment, because many VMs are using host network and CPU resources concurrently. As a result, backups that once completed during nonbusiness hours have seen their backup window grow.

NetApp SnapManager® for Hyper-V (SMHV) addresses the resource utilization issues typically found within virtual environments by leveraging the underlying NetApp Snapshot™ technology, thereby reducing the CPU and network load on the host platforms and drastically reducing the time required for backups to complete. SMHV can be quickly installed and configured for use in Hyper-V environments, saving valuable time during backups and allowing quick and efficient restorations, thus reducing administrative overhead.

1.1 Purpose and Scope

The purpose of the following chapters is to provide best practices for deploying SMHV to back up and recover Hyper-V VMs. They describe the key features and best practices to effectively manage the complete backup lifecycle for Hyper-V VMs. The present document only covers SnapManager for Hyper-V best practices in 7-Mode environments. For SnapManager for Hyper-V best practices on clustered Data ONTAP® 8.2, refer to TR-4226: NetApp SnapManager 2.0 for Hyper-V on Clustered Data ONTAP 8.2 Best Practices Guide.

1.2 Intended Audience

The following chapters are intended for Hyper-V administrators, storage administrators, backup administrators, and architects implementing a backup, restore, and disaster recovery solution for Hyper-V environments running on NetApp storage. Ideally, readers should have a solid understanding of the architecture, administration, and backup and recovery concepts within a Hyper-V environment and should consider reviewing the following documents:

- Data ONTAP 8.2 System Administration Guide
- SnapManager 2.0 for Hyper-V Installation and Administration Guide
- SnapDrive 7.0 for Windows Installation and Administration Guide

Technical Details

SMHV provides the following capabilities for Data ONTAP 8.2 operating in 7-Mode:

- Allows system administrators to create hardware-assisted backup and restore of Hyper-V VMs running on NetApp storage.
- Provides integration with Microsoft Hyper-V Volume Shadow Copy Service (VSS) writer to quiesce the Hyper-V VMs before creating an application-consistent Snapshot copy of the VM.
- Allows an administrator to create application-consistent backups of Hyper-V VMs, if the customer has Microsoft Exchange, Microsoft SQL Server®, or any other VSS-aware application running on virtual hard disks (VHDs) in the VM.
- Provides mirroring of backup sets to secondary locations for disaster recovery (DR) planning.
- Supports the backup and restore of shared VMs configured using Windows® failover clustering (WFC) for high availability and also on Microsoft cluster shared volumes (CSVs). SMHV makes sure that the scheduled VM backups can happen seamlessly irrespective of any VM failovers.

- Supports management of multiple remote Hyper-V parent systems from one console.
- Supports performing crash-consistent backup and restore of virtual machines.

2 SnapManager for Hyper-V Planning

2.1 Storage Considerations

SMHV supports backup and restore on CSVs. SMHV can back up only VM data stored in VHDs that reside on NetApp storage. It does not back up data on pass-through or direct-attached iSCSI disks. SMHV does not support MBR LUNs for VMs running on shared volumes or CSVs. SMHV 2.0 also supports backup of VMs in SMB 3.0 shares. However, this is only supported when connected to systems running clustered Data ONTAP 8.2 systems. SMHV supports LUNs created on thin-provisioned volumes and can perform backups/restores on these volumes.

3 SnapManager for Hyper-V Simplified Backup and Recovery

3.1 Prerequisites

SMHV supports backup and restore of virtual machines on dedicated disks, CSVs, and SMB 3.0 shares. SMHV can back up only VM data stored in VHDs that reside on NetApp storage. It does not back up data on pass-through or direct-attached iSCSI or VFC disks. SMHV does not support master boot record (MBR) LUNs for VMs running on shared volumes or CSVs. SMHV supports LUNs created on thin-provisioned volumes and can perform backups and restores on these volumes.

Note: SnapDrive® 7.0 for Windows (SDW) is required to be installed on the host systems. Using SnapDrive to provision LUNs or shares to host virtual machines is recommended.

Note: Backup and restore of VMs in SMB 3.0 shares are only supported with clustered Data ONTAP 8.2.

Note: SnapVault® updates are only supported with clustered Data ONTAP 8.2.

License Requirements

A SnapManager for Hyper-V license is required on the Windows host system. You can choose either host-based licensing or storage system licensing.

- If you select host-based licensing, you need to provide a license key during installation. You can change the license key after installation by clicking License settings in the SnapManager for Hyper-V Welcome window.
- If you select storage system licensing, you must add the SMHV license to all storage systems.

In addition, the following licenses are required:

- SnapRestore®
- The required protocol (FCP, iSCSI)
- SnapMirror® (if required)
- SnapDrive for Windows

Platform Support

- Windows Server® 2012 x64 Standard and Datacenter Editions (full and core installation)
- Hyper-V Server 2012 x64
- Windows Server 2008 x64 Standard, Datacenter, Enterprise, Editions (full and core installation)

- Hyper-V Server 2008 R2 SP1 x64

Remote Management Platform Support

- Windows Server 2012 x64 Standard, Datacenter (full installation)
- Hyper-V Server 2012 R2 x64 (full and core installation)
- Windows Server 2008 x64 Standard, Enterprise (full installation)
- Windows Server 2008 x64 Standard, Enterprise with SP2 (full installation)
- Windows Server 2008 R2 x64 Standard, Enterprise (full installation)
- Hyper-V Server 2008 R2 x64 (full and core installation)
- Windows Vista® x64 SP1, Windows Vista x86 SP1 and later
- Windows XP x86 with SP3 and later
- Windows Server 2003 x64 and x86 with SP2 and later
- Windows 8 and Windows 7

VM Support

- W2012 x64 Standard and Datacenter Editions (full and core installation)
- W2008 R2 x64 Standard, Datacenter, and Enterprise Editions (full and core installation)
- W2008 x64 Standard and Enterprise Editions (Full and Core)
- W2008 x64 Standard and Enterprise Editions with SP2 (full and core installation)
- W2003 x64 Standard and Enterprise Editions with SP2
- W2003 x86 Standard and Enterprise Editions with SP2
- Windows 8, Windows 7, Windows Vista, and Windows XP
- SUSE Linux® VMs (SLES10 SP-1 and SP2 - x86 and x64)
- RHEL 5.3, RHEL 5.4, and RHEL 5.5 (Microsoft Hyper-V Integration component version 2.1 must be installed)

For the most current information, refer to the NetApp Interoperability Matrix Tool.

Recommended Hotfixes

Table 1 lists the recommended hotfixes and updates.

Table 1) Microsoft hotfixes/updates.

Operating System	Hotfix
Windows Server 2012	2770917
	2779768
	2795944
	2811660
	2822241
	2836988
	2845533
	2851998
	2870270
Windows Server 2008 R2	978157
	979743

Operating System	Hotfix
	2406705
	974909
	975354
	977096
	974930
	2517329
	2779768
Windows Server 2008 SP1	2406705
	2531907
	2263829
	2494016
	2637197
	2517329
	2779768
	2494162

For the entire list of applicable KB articles, see the References section of this document.

4 SnapManager for Hyper-V Architecture

Table 2 lists the terminologies used throughout this document.

Table 2) Terminology.

Term	Description
Dataset	A dataset is a grouping of virtual machines that helps you to protect data using retention, scheduling, and replication policies. You can use datasets to group VMs that have the same protection requirements. A VM could be a member of multiple datasets. This can be useful for VMs that belong to multiple groupings (for example, a VM running the SQL Server instance for a Microsoft Office SharePoint® Server [MOSS] configuration might need to belong to both the SQL Server and the MOSS datasets).
Protection policies	Policies allow customers to schedule/automate the backups of the datasets at a predefined time (schedule policy), allow customers to provide retention capabilities for older backups (retention policy), and allow customers to replicate the block changes to the SnapMirror destination volume after the VM backup is created (replication policy). Policy includes other capabilities that allow customers to run scripts before and after the backup.
Backup and recovery	SMHV provides local backup and recovery capability with the option to replicate backups to a remote storage system using SnapMirror relationships. Backups are performed on the whole dataset, which is a logical collection of VMs, with the option of updating the SnapMirror relationship as part of the backup on a per-job basis. Similarly, restores can be performed at an individual VM level.
Application-consistent backup/restore	These backups are taken in coordination with the VSS to make sure that the applications running in the VM are quiesced before creating a Snapshot copy. Such a backup guarantees the integrity of application data and hence can be safely used to restore the VM and the applications running in the VM to a consistent state.

Term	Description
Crash-consistent backup	A backup in which the state of data is equivalent to what would be found following a catastrophic failure that abruptly shuts down the system. The data in the backup will be the same as it would be after a system failure or power outage. This type of backup is much quicker. A restore from such a backup would be equivalent to a reboot following an abrupt shutdown. **Note:** Crash-consistent backup and restore are supported from SMHV 1.1 onward and will require SnapDrive for Windows 6.4.1 to be installed on the host system.
Backup retention policy	Retention policies can be used to specify how long you want to keep a dataset backup based on either time or number of backups. Policies can be created specifying the retention period, allowing administrators flexibility to meet varying service-level agreement (SLA) levels within their environment.
Alert notification	Alert notifications are created on a per-scheduled-backup-job basis and are sent by e-mail to administrator-defined accounts. Alert notification can be configured to e-mail the specified account after every backup, although this is not recommended because the number of e-mails can become unmanageable. Configuring alerts to notify administrators after an error or warning within a backup offers a more useful and practical alert level.
Unprotected resources	Unprotected resources are VMs that are not part of any dataset. These resources can be protected by adding them to a dataset.

4.1 SnapManager for Hyper-V Port Usage

Best Practice

For SMHV and SDW, make sure that the following ports are kept open:

- 808: SMHV and SDW default port
- 4094: If SDW is configured to use HTTP protocol
- 4095: If SDW is configured to use HTTPS protocol

When SMHV is installed on a cluster, the same port number must be used across all nodes.

4.2 SnapManager for Hyper-V Architecture

SMHV must be installed on the Hyper-V parent host to create the backup and restore of the VMs running in the Hyper-V parent host. It provides an MMC-based management console and Windows PowerShell® snap-in for management operations. SMHV allows managing multiple Hyper-V parent hosts from a single console. The console can be installed on a Hyper-V parent host, other Windows (not Hyper-V) servers, and client systems such as Windows 8.

SMHV implements a VSS requestor, which communicates with the VSS framework and coordinates the application-consistent backups of the virtual machines. The administrator creates the dataset, composed of VMs from multiple physical hosts. After the dataset is created, various policies can be applied, composed of backup retention, backup schedule, and backup replication parameters. A replication policy provides the option to update SnapMirror and SnapVault.

When SMHV requests a backup or restore, the call is passed to SnapDrive for Windows, which is the VSS hardware provider. SnapDrive then communicates with Data ONTAP to create a hardware Snapshot copy.

Figure 1 **Error! Reference source not found.**illustrates the SMHV architecture and the components that work together to provide a comprehensive and powerful backup and recovery solution for Hyper-V environments.

Figure 1) SMHV architecture.

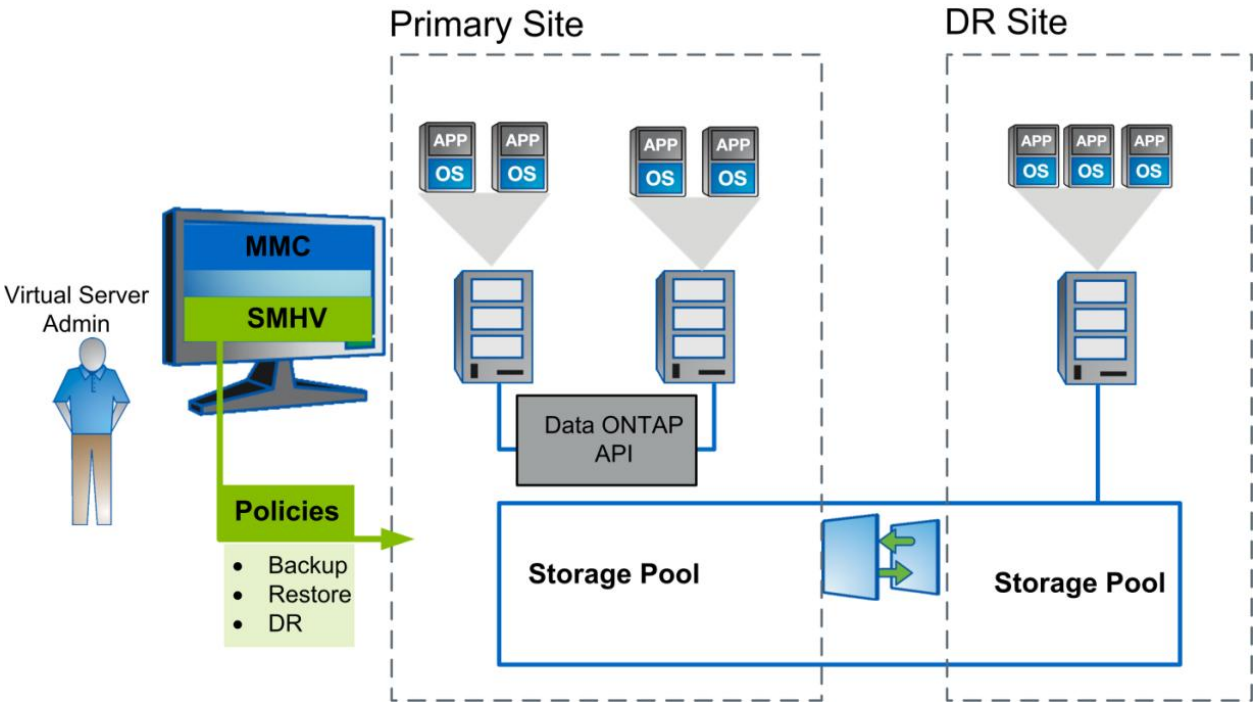

Components

SMHV SnapInfo Settings

The SMHV SnapInfo directory stores backup metadata. This folder can be set up by specifying the SnapInfo settings in the Hosts Management Wizard. The metadata information is critical to recovering VMs if a failure occurs. SnapInfo settings should be configured for the host or cluster added to SMHV so that VMs within that host can be added to a dataset. SMHV also creates a Snapshot copy of the SnapInfo directory after the backup is completed. The naming convention for the SnapInfo Snapshot copy is `smhv_snapinfo_hostname_timestamp`.

With SnapManager 2.0 for Hyper-V, SnapInfo can be hosted on CSV LUNs, SMB 3.0 shares, or dedicated LUNs.

Note: SnapInfo can be hosted on SMB 3.0 only for clustered Data ONTAP systems, and not for 7-Mode systems.

Note: If SnapInfo settings are changed, all files must be moved manually from the original SnapInfo location to the new location. SMHV does not move them automatically.

Best Practice
NetApp recommends having the SnapInfo LUN on a volume of its own.

SMHV Report Settings

Report settings should be configured for a host or cluster added to SMHV so that VMs within that host can be added to a dataset.

Best Practice
The report path must not reside on a CSV.

SMHV Event Notifications

Event notifications settings can be configured to send e-mail and AutoSupport™ messages in case an event occurs.

5 SnapManager for Hyper-V Backup Types

5.1 Application-Consistent Backup

Microsoft Volume Shadow Copy Service was developed specifically to enable third-party backup and recovery solutions to provide application-consistent backup and recovery for mission-critical applications supported by Microsoft. When VSS is properly configured in the Hyper-V environment, an SMHV-initiated Snapshot copy begins the VSS process.

VSS is designed to produce fast, consistent Snapshot copy-based online backups by coordinating backup and restore operations among business applications, file system services, backup applications, fast-recovery solutions, and storage hardware.

VSS coordinates Snapshot copy-based backup and restore and includes these additional components:

- **VSS requestor.** The VSS requestor is a backup application, such as the SMHV application or NTBackup. It initiates VSS backup and restore operations. The requestor also specifies Snapshot copy attributes for the backups it initiates.
- **VSS writer.** The VSS writer owns and manages the data to be captured in the Snapshot copy. Microsoft Hyper-V is an example of a VSS writer.
- **VSS provider.** The VSS provider is responsible for creation and management of the Snapshot copy. A provider can be either a hardware provider or a software provider: A hardware provider integrates storage array–specific Snapshot copy and cloning functionality into the VSS framework. The Data ONTAP VSS hardware provider integrates the SnapDrive service and storage systems running Data ONTAP into the VSS framework. A software provider implements Snapshot copy or cloning functionality in software running on the Windows system.

The coordinated backup process includes:

- Freezing the data application I/O
- Flushing the file system cached I/O to disk
- Creating a point-in-time Snapshot copy of the data state

After the Snapshot copy is created, file system and application I/O resume. The VSS restore process involves:

- Placing the data application into the restore state
- Passing backup metadata back to the application whose data is being restored
- Restoring the actual data
- Signaling the data application to proceed with recovering the data that was restored

SMHV provides integration with Microsoft Hyper-V VSS writer to quiesce a VM before creating an application-consistent Snapshot copy of the VM. SMHV is a VSS requestor and coordinates the backup operation to create a consistent Snapshot copy, using VSS hardware provider for Data ONTAP. SMHV makes it possible to create application-consistent backups of a VM if Microsoft Exchange, Microsoft SQL Server, or any other VSS-aware application is running on VHDs in the VM. The applications that exist in the VM are restored to the same state that existed at the time of the backup. SMHV restores the VM to its original location.

If applications are running on pass-through or direct-attached iSCSI LUNs, these LUNs are ignored by the VSS framework in the VM, and SMHV does not create a backup of these LUNs in the VM. To enable backup of application data on direct-attached iSCSI LUNs or pass-through LUNs in the VM, it is necessary to configure application backup products in the VM (for example, SnapManager for Exchange, SnapManager for SQL Server, and so on).

Note: The Data ONTAP VSS hardware provider is installed automatically as part of the SnapDrive software installation.

To make sure that the Data ONTAP VSS hardware provider works properly, do not use the VSS software provider on Data ONTAP LUNs. If the VSS software provider is used to create Snapshot copies on a Data ONTAP LUN, that LUN cannot be deleted by using the VSS hardware provider.

Note: VSS requires the provider to initiate a Snapshot copy within 10 seconds. If this time limit is exceeded, the Data ONTAP VSS hardware provider logs event ID 4364. This limit could be exceeded due to a transient problem. If this event is logged for a failed backup, retry the backup.

Note: SMHV coordinates with Hyper-V VSS writer to create application-consistent backup of VMs. Hyper-V writer communicates with integration services (Hyper-V VSS requestor service) installed in the VM to quiesce the applications running in the VM before creating a backup. Data ONTAP VSS hardware provider installed on the Hyper-V host as part of SnapDrive is used to create Snapshot copies on the storage system.

For detailed information about VM backup, refer to Planning for Backup on the Microsoft TechNet site.

5.2 Crash-Consistent Backup and Restore

Backups created using SMHV can be either application-consistent or crash-consistent. Application-consistent backups are created in coordination with Microsoft Volume Shadow Copy Service (VSS) to make sure that the applications running in the VM are quiesced before creating the Snapshot copy. Such a backup assures the integrity of application data; therefore, it can be safely used to restore the VM and the applications running in the VM to a consistent state.

Although application-consistent backups are the most suitable solution for data protection and recovery of Hyper-V VMs, they also have a few drawbacks:

- Application-consistent backups are slower due to VSS involvement with the parent and guest OS. Because both the application writer in the VM and the Hyper-V writer in the parent OS are involved, failure to back up any of the components will cause the backup process to fail.
- Hyper-V writer uses the autorecovery process to make the VMs consistent. Autorecovery results in the creation of two Snapshot copies on the storage system. Therefore, each Hyper-V backup requires two Snapshot copies to be created per storage system volume.
- If multiple VMs are running on different nodes in a cluster, but on the same CSV, SMHV still needs to create one backup per node as required by VSS. As a result, SMHV creates multiple Snapshot copies on the same CSV for different VMs.

Considering these drawbacks, it is desirable to have some way of creating "quick" Hyper-V VM backups. Crash-consistent backup is designed to provide this ability of creating quick backups.

A crash-consistent backup of a VM does not use VSS to quiesce data, and it does not result in autorecovery. This backup simply creates a Snapshot copy on the NetApp storage system for all the

LUNs used by the VMs involved in the dataset. The data in the backup is the same as it would be after a system failure or power outage. All of the SMHV functions such as scheduling, restore, script execution, SnapMirror updates, backup retention, and so on are supported for crash-consistent backups as well.

Figure 2 illustrates the Backup Dataset Wizard showing the application-consistent and crash-consistent backup types.

Figure 2) Backup Dataset Wizard screen.

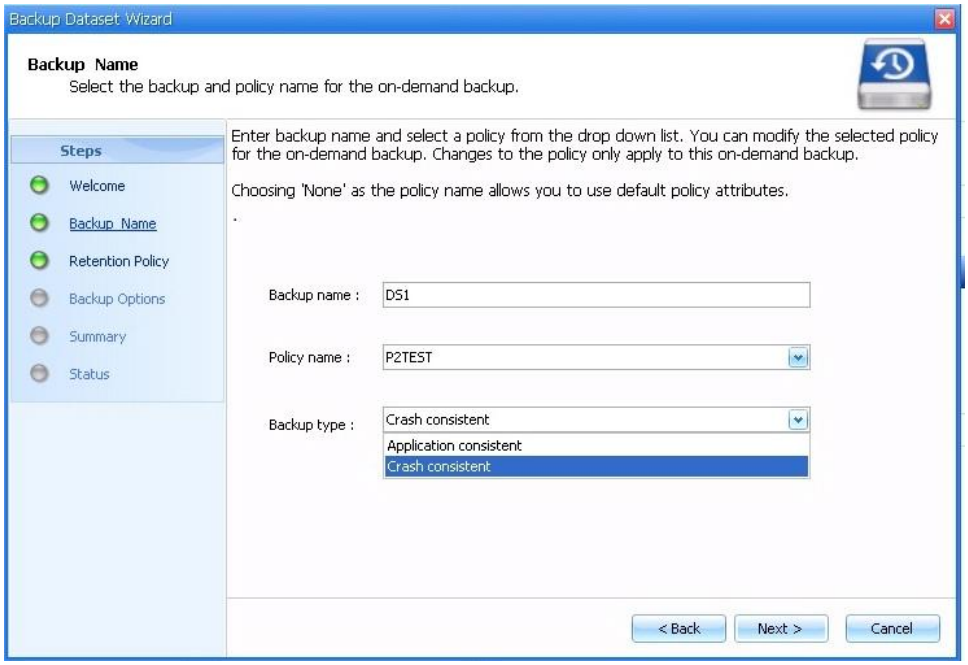

Note: Saved state backup policy is not applicable for crash-consistent backup and restore. This is because crash-consistent backups do not involve the Hyper-V VSS writer.

Note: SMHV supports parallel execution crash-consistent and application-consistent backups. It also supports parallel crash-consistent backup execution. However, users might observe some issues while such operations are executed. This is due to a timeout error in the underlying SnapDrive for Windows.

Best Practice

The crash-consistent backup feature is not a replacement for application-consistent backups. It enables the user to have frequent recovery points and therefore to have frequent crash-consistent backups and fewer application-consistent backups.

Best Practice

Crash-consistent backup can be used to create the latest backup of all the data just before performing an application-consistent restore operation of a VM.

6 SnapManager for Hyper-V Process Flow

6.1 SnapManager for Hyper-V Installation

As discussed previously, SDW 7.0 is a prerequisite for SMHV 2.0 installation. Both SDW and SMHV must be installed on all the nodes in case of a Windows clustered environment. SMHV 2.0 supports the remote installation of SMHV from one server to another server. It can be used to remotely install across all the partner nodes in a Windows cluster. You can install SMHV remotely by providing the host name and credentials of the remote server along with SMHV and SDW licenses.

If the SMHV and SDW licenses are already installed on the storage system, the “per storage” option can be selected.

Note: Remote installation is supported for standalone and cluster nodes within a domain. SMHV cannot be installed on a host, which is part of another domain.

6.2 Adding a Hyper-V Parent Host or Host Cluster

If a single host is added, SMHV manages the dedicated VMs on that host. If a host cluster is added, SMHV manages the shared VMs on the host cluster. If there is a plan to add a host cluster, SMHV must be installed on each cluster node. SMHV 2.0 supports remote installation of SMHV on all nodes of the Windows cluster from a single node.

If the backup repository settings, report directory settings, and notification settings are not configured for SMHV, they can be configured after the host is added by using the configuration wizard. The backup repository and report directory settings must be configured in order to add and manage VMs using SMHV. Notification settings are optional.

Note: Dedicated and shared VMs that belong to the same host cluster should not exist in the same dataset. Adding these types of resources to a single dataset can cause the dataset backup to fail.

Although a host should be managed from only one management console, if the need arises, it is possible to do so from multiple consoles. It is possible to import and export host and dataset configuration information from one remote management console to another for data consistency. The Import and Export Wizard can also be used to change host and dataset configuration settings to previously exported settings. If this operation is performed in a clustered environment, the settings on all nodes in the cluster must be imported so that all host and dataset configurations are the same.

Caution

Do not import or export configuration information to the directory where SMHV is installed, because if SMHV is uninstalled this file will be lost.

6.3 SMHV Backup Mechanism in Windows Server 2008 R2 SAN Environments

SMHV leverages NetApp Snapshot technology to create fast and space-efficient backups of SMHV datasets and their associated VMs. These backups offer point-in-time images, or copies, of the VMs and are stored locally on the same storage platform on which the VMs reside.

In addition to the Snapshot copy stored locally, SMHV also provides an option to update an existing SnapMirror or SnapVault relationship at the completion of a backup. The administrator can select this on a per-backup-job basis. The unit of backup in SMHV is a dataset, which can contain one or more VMs running across multiple Hyper-V hosts. SMHV supports restoring an individual VM; it does not support restoring an entire dataset. Using SMHV, on-demand or scheduled backups of VMs are possible. SMHV supports backup of dedicated or clustered VMs. It also supports backups of shared VMs running on CSVs and SMB 3.0 shares (clustered Data ONTAP 8.2 only).

Figure 3 illustrates a high-level overview of the typical SMHV architecture on the primary site storage, where the backup process takes place in a SAN environment.

Figure 3) Hyper-V infrastructure and associated storage during an application-consistent SMHV backup.

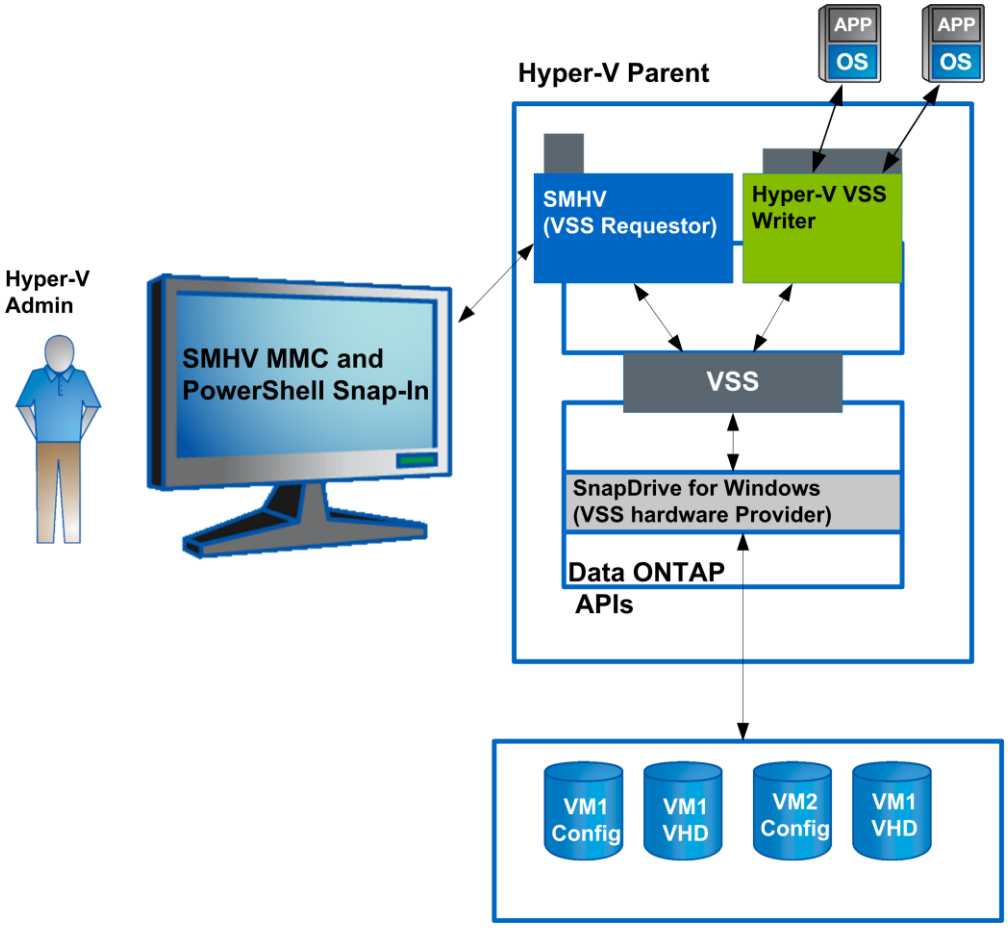

The following steps describe the flow of the backup process in a SAN environment:

1. The SMHV service, a VSS requestor, initiates a VSS backup of VMs within a dataset in coordination with the Microsoft Hyper-V VSS writer.
2. The Hyper-V VSS writer works together with the integration services within the VM to create application-consistent software Snapshot copies of all VHD volumes attached to each VM.
3. SMHV implements a VSS requestor component to coordinate the backup process and create a consistent Snapshot copy in Data ONTAP using VSS hardware provider for Data ONTAP LUNs.
4. The VSS framework asks the hardware provider to mount the LUNs from the Snapshot copy.
5. The Hyper-V writer recovers data on the LUNs and brings it to the state of the software Snapshot copy that was created in step 2.
6. The VSS provider creates a second Snapshot copy of the LUNs and then dismounts them from the Snapshot copy.
7. At the completion of the local backup, SMHV updates an existing SnapMirror relationship on the volume if the Update SnapMirror option was selected.

SMHV makes it possible to create application-consistent backups of a VM if Microsoft Exchange, Microsoft SQL Server, or any other VSS-aware application is running on VHDs in the VM. SMHV

coordinates with the application VSS writers inside the VM so that application data is consistent when the backup occurs.

Note: For a backup to succeed, all files of the VM (VHDs, VM configuration files, and VM Snapshot files) should reside on LUNs managed by Data ONTAP.

Note: Only one backup operation can occur on a host at any given time. If the same VMs belong to different datasets, do not schedule a backup of the datasets at the same time. If this occurs, one of the backup operations will fail.

Note: SMHV backup fails for VMs that have a VHD created by copying the contents of a physical disk on the same host. The Create New VHD wizard of Hyper-V Manager provides this option. As part of copying the physical disk contents, it also copies the disk signature, which causes disk signature conflict during the backup. More information is available at the Microsoft Support web site.

Note: SnapVault updates in SMHV 2.0 are not supported with Data ONTAP 7-Mode systems.

Caution

Do not create a VHD by using the option Copy the Contents of the Specified Physical Disk in the Configure Disk page in the Create New VHD Wizard in Microsoft Hyper-V Manager.

Note: SMHV does not support the backup and restore of VMs running on SAN boot LUNs.

Note: Grouping of virtual machines hosted on SMB shares and SAN LUNs in a single dataset is not supported.

Workflow for crash-consistent backups:

1. The user chooses the crash-consistent backup option in the Backup Dataset Wizard.
2. The SMHV API calls VSS to collect the VM metadata. The LUNs on which the VMs are hosted are identified.
3. The SnapDrive API is called to create a Snapshot copy of these LUNs. Only one Snapshot copy is created for each LUN, regardless of the number of VMs running on it.
4. The backup is registered with the backup type Crash-Consistent.
5. Upon completion of the local backup, if the SnapMirror option is selected, SMHV updates an existing SnapMirror relationship on the volume.

Note: While performing a crash-consistent backup or restore, SMHV does not leverage VSS. VSS is used only to get VM-related metadata from the Hyper-V writer. The default backup type is Application-Consistent.

Best Practice

When creating a dataset, select all VMs that reside on a particular Data ONTAP LUN. This makes it possible to get all backups in one Snapshot copy and to reduce the space consumption on the storage system. It is preferable to add VMs running on the same CSV in the same dataset. If VMs are added on the same CSV in different datasets, make sure that the backup schedules of these datasets do not overlap.

Best Practice

If a VM Snapshot copy location is changed to a different Data ONTAP LUN after the VM is created, create at least one VM Snapshot copy by using Hyper-V Manager before creating a backup by using SMHV. If this is not done, the backup could fail.

6.4 SMHV Backup Mechanism for Windows Server 2012 SAN Environments

In Windows Server 2012, Microsoft introduced the **CSV proxy file system (CSVFS), which** provides a cluster shared storage LUN with a single and consistent file namespace while still using the underlying NTFS file system. In Windows Server 2012, the CSVs now appear as CSV file system, instead of NTFS (as in Windows Server 2008 R2). For additional information on CSVFS architecture, refer to Introduction to Cluster Shared Volumes and CSV Architecture.

In Windows Server 2008 R2, CSV Hyper-V backup creates application-consistent backups on each VM owner node. CSV ownership is moved to the VM owner node as part of the backup process. Hyper-V VSS writer then coordinates the freeze and thaw operations in the Hyper-V guest, and a subsequent hardware Snapshot copy is created from the Hyper-V parent using the Data ONTAP VSS hardware provider (SnapDrive for Windows). This results in the creation of a hardware Snapshot copy for each Windows cluster node, thereby introducing several scalability and space efficiency issues when the number of nodes in the cluster increases.

In Windows Server 2012, CSVFS introduced “distributed application-consistent backups.” This allows backup of all the VMs in a cluster to be consistent in one single application-consistent backup. To achieve this distributed backup mechanism, Microsoft has introduced a new CSV writer and CSV provider.

- **CSV writer.** CSV writer serves the component-level metadata from the nonrequesting node for CSV volumes, and it functions as a proxy by including the Hyper-V writers from the remote node for the backup session.
- **CSV provider.** CSV provider coordinates the VSS backup activities from all the Hyper-V writers on the partner cluster nodes to make the VM in an application-consistent state. Also, the CSV provider makes sure that CSV shadow copy volume is writable for the partner node Hyper-V writers during the autorecovery process.

Figure 4 illustrates the SMHV backup process for Windows Server 2012.

Figure 4) SMHV backup process for Windows Server 2012 SAN environments.

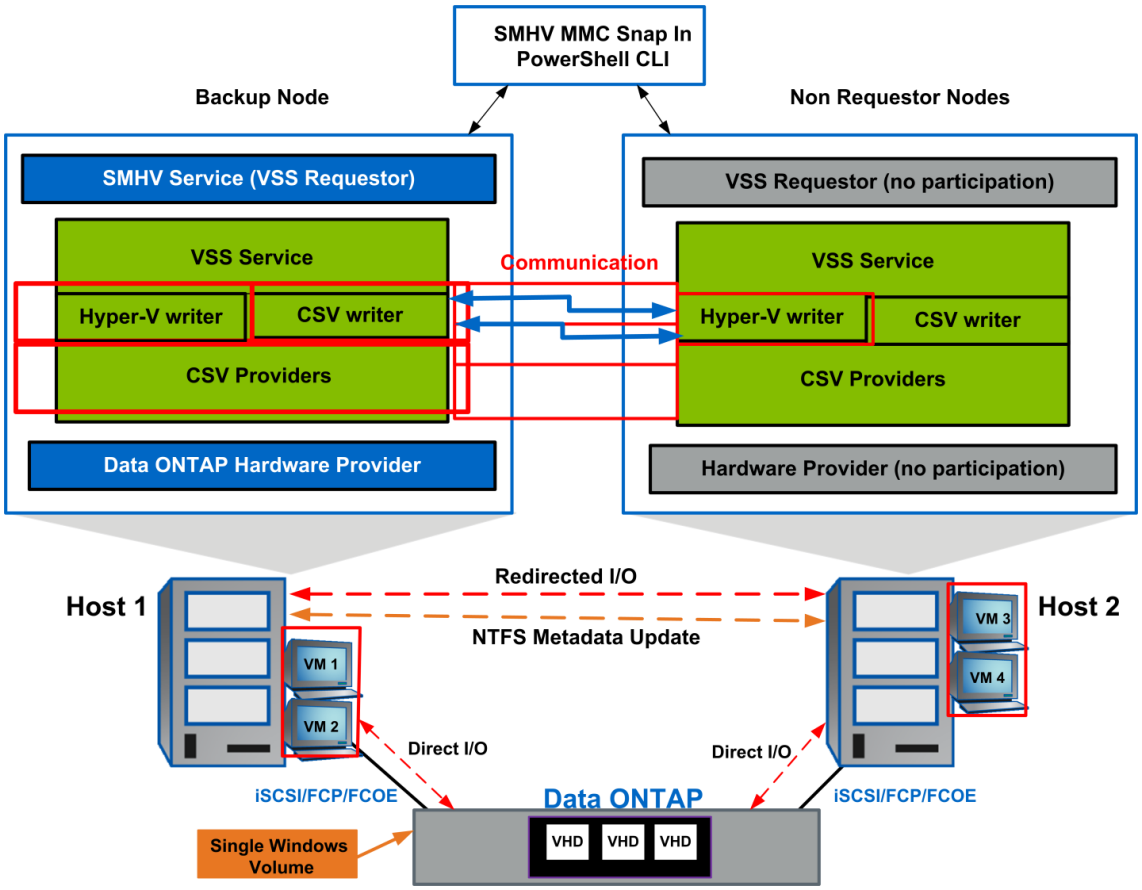

Initialization Phase

- The user initiates the backup operation from any node in the cluster by using SMHV. SMHV redirects the backup operation to the Windows cluster owner node, which functions as a coordinator node throughout the entire backup operation.
- SMHV initializes the Microsoft VSS operation only in the coordinator node. This is unlike Windows Server 2008 R2, in which VSS is initialized in each node of the Windows cluster that is involved in the backup. This optimization improves the overall timing of the backup operation.
- SMHV gathers the metadata (files used by VMs) for all the VMs involved in the backup. Metadata for VMs that are local to the coordinator node is gathered by the Hyper-V writer running in the coordinator node.
- Metadata for the VMs that are not local to the coordinator node is gathered by the CSV writer running in the coordinator node. Internally the CSV writer in the coordinator node interacts with the Hyper-V writer in other nodes to get the metadata from all other nodes. So, unlike Windows Server 2008 R2, in which SMHV explicitly reaches out to each node to capture the metadata, this complication is handled by the new CSV writer in Windows Server 2012.

Prebackup Phase

- The Hyper-V writer on the coordinator node quiesces the application writers inside the VM by using the integration service.
- The CSV software provider on the coordinator node interacts with the Hyper-V writers in all the other VM owner nodes to make sure that the state of the application running inside the VM is consistent before starting the actual Snapshot copy of the volume.

Backup Phase

- The VSS hardware provider on the coordinator node creates the backup Snapshot copy of the CSV volume.
- After the hardware Snapshot copy is created, the Hyper-V writer, by default, performs an autorecovery process on each VM owner node to remove any in-flight transactions. Autorecovered changes are applied on the pseudo CSV Snapshot disk object exposed on the backup node, which is accessible from all the other VM nodes. This process makes the backups on each VM owner node application consistent with respect to the CSV.

Postbackup Phase

- SMHV retrieves the VSS backup metadata and backup component documents and then modifies the metadata to make it compatible with VSS-required semantics.
- SMHV saves the backup metadata to the SnapInfo directory.
- The VSS Snapshot GUID is renamed to the SMHV naming conventions.
- Applicable policy processing such as retention of older backups, SnapMirror updates, running any specified postscript, or generating ASUP™ notifications, is performed.

Note: Make sure that the Enable Distributed Backup option is selected in the Backup Dataset Wizard.

Note: The distributed backup mechanism for Windows Server 2012 is not applicable for the crash-consistent backup feature in SMHV.

Note: NetApp recommends that all the VHD files belonging to a virtual machine should be hosted on CSVFS LUNs only, not on a mix of CSVFS and shared disks. This is because SMHV does not support such mixed-mode backups.

To summarize, distributed application-consistent backups are faster because they avoid multiple backup requests to each node in the cluster. The entire backup operation is performed from the coordinator node (cluster owner) alone and by leveraging the new CSV writer and CSV shadow copy provider.

Also, distributed application-consistent backup is more space efficient because it creates only one Snapshot copy for each volume instead of creating one Snapshot copy for each node and volume combination. This space saving is huge if large numbers of nodes are involved in the backup. Also, Data ONTAP imposes a limit for the maximum number of Snapshot copies that can be stored for a volume, so this enhancement allows storing more backups for a VM.

Note: SnapManager for Hyper-V does not support having virtual machines in such CSVs hosted on asymmetric clusters in Windows Server 2012.

Note: SnapManager for Hyper-V supports BitLocker functionality for CSVs provisioned through SnapDrive for Windows. Virtual machines (VMs) can be hosted in encrypted CSVs in Windows Server 2012.

Note: SnapDrive for Windows currently cannot create LUNs beyond 14TB, and therefore NetApp recommends creating a VHDx for sizes less than 14TB and to use other means of provisioning additional storage (pass-through disks, guest iSCSI initiator) on the VM.

Note: SMHV supports backup of VMs that have LUNs attached with virtual Fibre Channel in the VM. However, during the VM backup, the LUN presented to the VM is not backed up.

Note: In Windows Server 2012, users can perform concurrent live migration of multiple VMs from one node to another. It is best to avoid SMHV-related operations within the virtual machine during live migration.

Note: Windows Server 2012 enables migrating virtual machine-r-related files to a different storage location without the VM having to undergo downtime. It is no longer necessary to take the virtual machine state offline when migrating to a different storage system. After migrating the virtual machine from one volume to another, restoring to a Snapshot copy taken in the earlier volume is not supported. It is best to avoid SMHV-related operations during storage live migration. Otherwise, such operations could corrupt the virtual machine.

6.5 SMHV Backup Mechanism for Windows Server 2012 SMB 3.0 Environments

Data ONTAP 8.2 operating in 7-Mode does not support Hyper-V over SMB. Hosting of VMs in SMB 3.0 shares is only supported with clustered Data ONTAP 8.2. For understanding SMHV backup architecture in Hyper-V over SMB environments, refer to TR-4226: NetApp SnapManager 2.0 for Hyper-V on Clustered Data ONTAP 8.2 Best Practices Guide.

6.6 Scheduled Backups and Retention Policies

SMHV allows administrators to schedule a dataset backup at a particular time. SMHV uses the Windows Tasks Scheduler for creating or modifying scheduling policies. The 255 NetApp Snapshot copies-per-volume limit must be taken into consideration when scheduling backups and configuring the associated retention policies. The number of Snapshot copies per volume can be managed with the proper scheduling and retention policies on a per-scheduled-backup basis while still meeting SLAs on the VMs.

Backup Scheduling

Using scheduling policies, administrators can schedule backup jobs at particular times, allowing them to automate the process. Multiple policies can be scheduled per dataset that apply to all hosts that are dataset members.

Best Practice

The backup frequency, as well as the number of different backups performed against a dataset—for example, one backup running against dataset ds_1 weekly and another monthly—must be taken into account when specifying the retention policy so as not to exceed the maximum number of Snapshot copies per volume. Should the number of Snapshot copies exceed 255 on any given volume, future backups against that volume will fail.

Retention Policies

The following list describes the retention tags available in SMHV:

- **Hourly.** Hourly intervals
- **Daily.** A specified time within a 24-hour period
- **Weekly.** A specified day and time within a seven-day period
- **Monthly.** A specified day and time within a calendar month
- **Unlimited.** Never-deleted backups

After choosing a retention type, you can choose to delete either backups that are older than a specified period of time or backups that exceed a maximum total.

NetApp recommends using the policies not only to meet specific SLAs, but also to maintain a supported number of NetApp Snapshot copies on the underlying volumes. For SMHV, one backup creates two Snapshot copies on the storage systems for data consistency (refer to KB ID 2010607). For example, setting a retention policy of 30 backups on an hourly basis limits the maximum number of Snapshot copies associated with the backup to 60. However, if the retention policy had been configured as 30 days, the Snapshot limit per volume would be reached in 5 days, and backups would begin to fail from that point on.

Best Practice

Choose a backup retention level based on your backup creation and verification schedule. If a Snapshot copy deletion occurs, make sure that a minimum of one verified backup remains on the volume. Otherwise, you run a higher risk of not having a usable backup from which to restore in case of a disaster.

Note: The option "unlimited" should be used with caution. When this option is selected, backups and the associated NetApp Snapshot copies are maintained until they are manually deleted by the administrator. These Snapshot copies are included in the maximum number supported on a volume.

Of further note, the NetApp Snapshot copies associated with on-demand backups must also be considered when determining the number of Snapshot copies maintained against a volume.

After creating a dataset backup, SMHV creates a Snapshot copy of the SnapInfo LUN. SnapInfo Snapshot copies are not deleted if the backup is deleted. SnapInfo Snapshot copies have a different retention policy. By default, SMHV retains 30 SnapInfo LUN Snapshot copies and deletes the older ones when the SnapInfo Snapshot count exceeds 30. You can configure the number of SnapInfo Snapshot copies you want to retain for each Hyper-V host using the following registry key:

- For standalone Hyper-V hosts:
  Registry key: `HKLM\SOFTWARE\NetApp\SnapManager` for Hyper-V\Server DWORD value:
  `snapinfo_snaps_count` (number of SnapInfo Snapshot copies to be retained)
- For clustered Hyper-V hosts (to be configured on each node in the cluster):
  Registry key: `HKLM\Cluster\SOFTWARE\NetApp\SnapManager` for Hyper-V\Server DWORD value:
  snapinfo_snaps_count (number of SnapInfo Snapshot copies to be retained)

6.7 Handling Saved-State Backup of VMs

The default behavior of SMHV is to fail a backup if one or more VMs cannot be backed up online. If a VM is in the saved state or shut down, an online backup cannot be performed. In some cases, VMs are in the saved state or shut down for maintenance, but backups still need to proceed, even if an online backup is not possible. To do this, the VMs that are in the saved state or shut down can be moved to a different dataset with a policy that allows saved-state backups.

Note: You can also select the Allow saved-state VM backup checkbox to allow SMHV to back up the VM using the saved state. If you check this option, SMHV will not fail the backup when the Hyper-V VSS writer backs up the VM using the saved state or performs an offline backup of the VM. Doing a saved state or offline backup can cause downtime. For more information on online or offline VM backups, see the Hyper-V Planning for the Backup information in the Technet library: http://technet.microsoft.com/en-us/library/cc753637(WS.10).aspx.

Best Practice

For mission-critical VMs, NetApp recommends disabling the "Allow Saved state VM backup" option.

Note: The "Allow saved state policy" option is not applicable for crash-consistent backups. This is because the VM is being backed up irrespective of the state.

6.8 Backup Scripts

Using SMHV, you can run optional backup scripts either before or after the backup takes place. These scripts will run on all dataset member hosts unless you indicate a specific server. The following environment variables can be used as arguments for application-consistent backup postscripts:

- `$VMSnapshot`

  Specifies the first VM Snapshot copy name that is created on a storage system as a result of the backup. The second name uses the first name plus the `appendix _backup`.

- `$SnapInfoName`

  Specifies the time stamp used in the SnapInfo directory name.

- `$Snapinfosnapshot`

  Specifies the SnapInfo Snapshot copy name created on the storage system. SMHV makes a Snapshot copy of the SnapInfo LUN at the end of the dataset backup.

During the post–script execution phase, SMHV replaces the $VMSnapshot variable with the Snapshot name, `$SnapInfoName` with the time stamp of the backup, and `$SnapInfoSnapshot` with the SnapInfo Snapshot name.

Note: The $SnapInfoSnapshot variable is supported for dedicated virtual machines only.

6.9 Quick/Live Migration Implications

Best Practice

SMHV cannot back up a VM that is actively undergoing migration. When a backup runs against a dataset that has VMs actively being migrated, an error is generated, and those particular VMs are not backed up.

6.10 Restore Process

SMHV can restore a VM from a backup. SMHV can also restore a VM that is part of a cluster. To restore the VM, SMHV uses the file-level restore feature in SDW. You can spread the associated files of a VM, including the configuration file, Snapshot copies, and any VHDs, across multiple Data ONTAP LUNs. A LUN can contain files belonging to multiple VMs.

If a LUN contains only files associated with the VM you want to restore, SMHV restores the LUN using LUN clone split restore (LCSR). If a LUN contains files not associated with the VM you want to restore, SMHV restores the VM using the file copy restore operation.

With these differences in restore types aside, the process flow used by SMHV during a restore is as follows:

1. SMHV restores a VM in coordination with Hyper-V VSS writer. Hyper-V VSS writer powers off the VM and deletes it before restore.
2. Files are restored as described in the preceding paragraphs based on restore type.
3. SMHV notifies the VSS writer that the files of the VM are restored properly. Hyper-V VSS writer registers the VM, and the VM gets added back in the Hyper-V manager.
4. SMHV starts the VM after restore and executes a postscript if specified in the restore wizard.

Note: During the restore, the following warning messages might be displayed:

- VM to be restored is not [currently running] on the host.
- VM to be restored is currently running on the host, and:
  - It has more VHDs associated with it than at the time of backup.

  - It has fewer VHDs associated with it than at the time of backup.
- The Snapshot location of the VM has changed.
- The names of VHD files or their file system paths or NetApp storage system LUN paths have changed.

In all of these warning scenarios, the VM can be restored, but you must acknowledge that you are sure you want to go ahead with the restore.

Note: If the VM no longer exists, you can still restore it if the LUNs on which the VM was created still exist. The LUNs must have the same drive letters and Windows volume GUIDs as at the time of backup.

If the VM no longer exists, you can still restore it by selecting a backup to which it belonged.

If the VM was removed from all datasets before it was deleted, you can still restore it by selecting unprotected resources and selecting a backup to which it belonged.

Best Practice

If the number of VHDs attached to a VM at the time of backup and restore is not same, the restored VM might have additional/fewer VHDs. If that is the case, NetApp recommends that the cluster configuration of the VM and its dependencies be manually updated.

Note: SMHV does not back up the cluster configuration of the VM, so it does not restore the cluster configuration. If the VM and the cluster configuration are lost, you can restore the VM from SMHV, but you must manually make it highly available. For more information, see "Failover Clustering on Windows Server 2008 R2" on the Microsoft website.

Note: In case of crash-consistent backups, the VM is restored without involving the VSS. It performs a file-level restore of the VM using SnapDrive for Windows.

Note: Restoring a deleted VM is not supported for crash-consistent backups. Also, RestoreToAlternateHost switch in Restore-Backup cmdlet cannot be used when the backup being restored is a crash-consistent backup.

6.11 Mounting a Backup

Backups can be mounted using SnapDrive for Windows. The mounted backup is a clone of the protected VM. After being mounted, the backup is displayed within the explorer of Hyper-V host and can be browsed.

1. Select the LUN, and within Snapshot copies select the backup to mount.

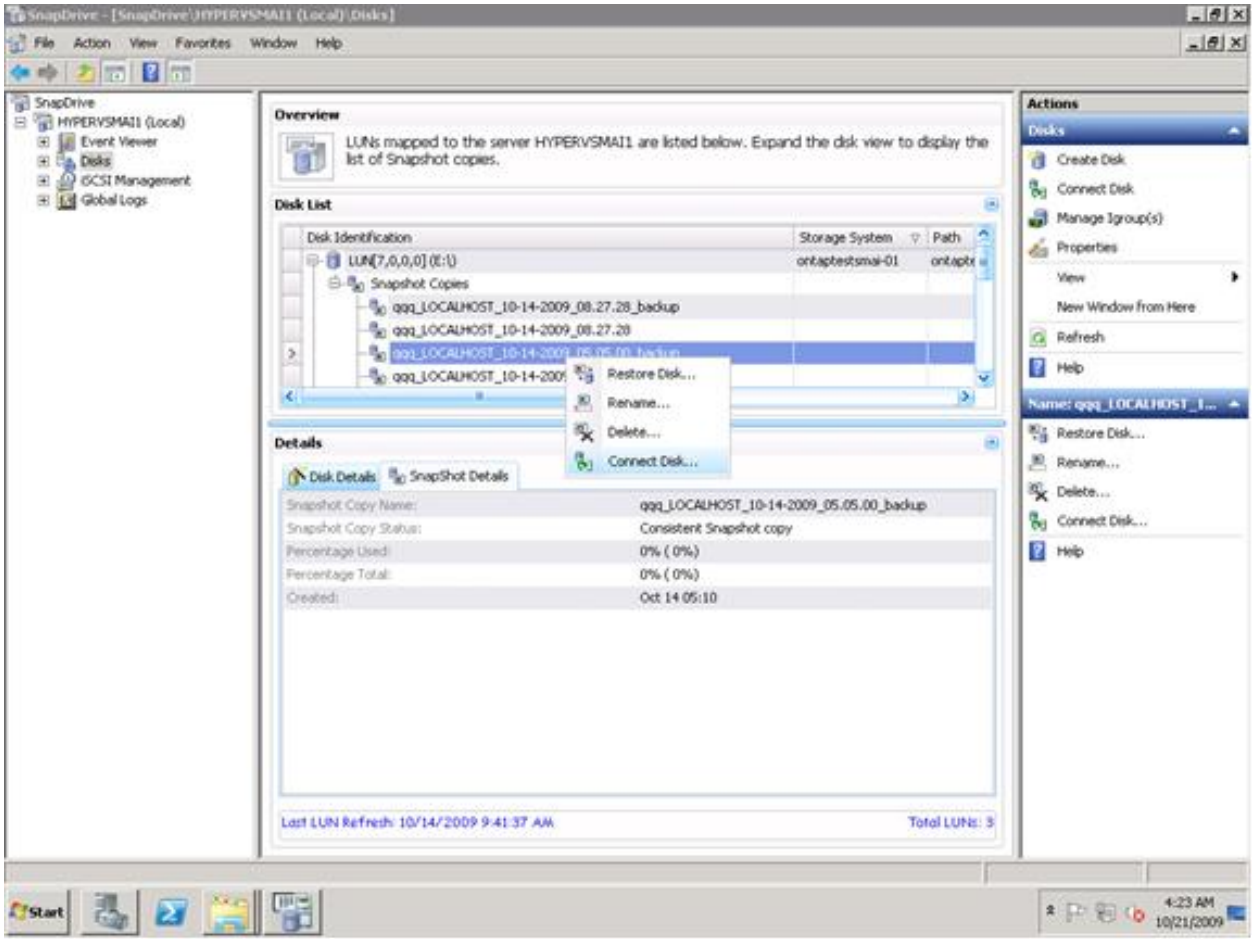

2. Right-click the Snapshot copy (the one with _backup suffix) and select the connect disk option.
3. Click Next.
4. If the LUN is a dedicated disk, go to the next step; otherwise, if the LUN is a Windows cluster resource, perform the following steps in the Specify Microsoft Cluster Services Group panel. In the Specify Microsoft Cluster Services Group panel, perform one of the following actions and then click Next.
    a. Select a cluster group from the Group Name drop-down list.
    b. Select Create a new cluster group to create a new cluster group.

Note: When selecting a cluster group for your LUNs, choose the cluster group your application will use.

Note: If you are creating a volume mount point, the cluster group is already selected. This is because the cluster group owns your root volume physical disk cluster resources. NetApp recommends that you create new shared LUNs outside of the cluster group.

    c. Select Add to cluster shared volumes.

5. In the Select LUN Properties panel, perform the following actions: Either select a drive from the list of available drive letters, or enter a mount point for the LUN you are connecting. When you create a volume mount point, enter the drive path that the mounted drive will use: for example, G:\mount_drive1\.
6. In the Select Initiators panel, choose an initiator for the LUN.
7. In the Select Initiator Group Management panel, specify whether you will use automatic or manual igroup management.
8. In the Completing the Connect Disk Wizard panel, perform the following actions:
   a. Verify all the settings.
   b. If you need to change any settings, click Back to go back to the previous wizard panels.
   c. Click Finish.
9. Browse the backup by selecting the drive letter on the explorer of Hyper-V host.

Single-File Restore Capability

In addition to backup verification, mounting a backup provides a way to restore a single file from within a VM on a case-by-case basis. This is performed by attaching a VHD from within the mounted backup as an existing hard drive to a VM within Hyper-V manager. After a backup has been mounted, the user can use the “Hot Disk Add” functionality in Windows 2008 R2 to attach a disk (backed by the VHD) to the VM at run time without shutting down the VM. Using this functionality, the user can attach new disks to the VM.

This is a three-step process, as the following procedure shows:

1. The user must first mount the VHD from the backup mounted location [<drive>:\ Name.vhd] to the parent host using the Attach VHD option from Disk Management SnapIn. This mounts the VHD as a new disk in the Hyper-V parent.
2. Offline the disk just mounted in the preceding step using the Disk Manager Snapin. Select the disk and choose the offline menu item. This offlines the disk mounted from VHD.
3. Attach the offlined disk to the virtual machine by selecting the Physical hard disk radio button and choose the disk that just offlined from the VM settings property page, as shown in the following screen capture.

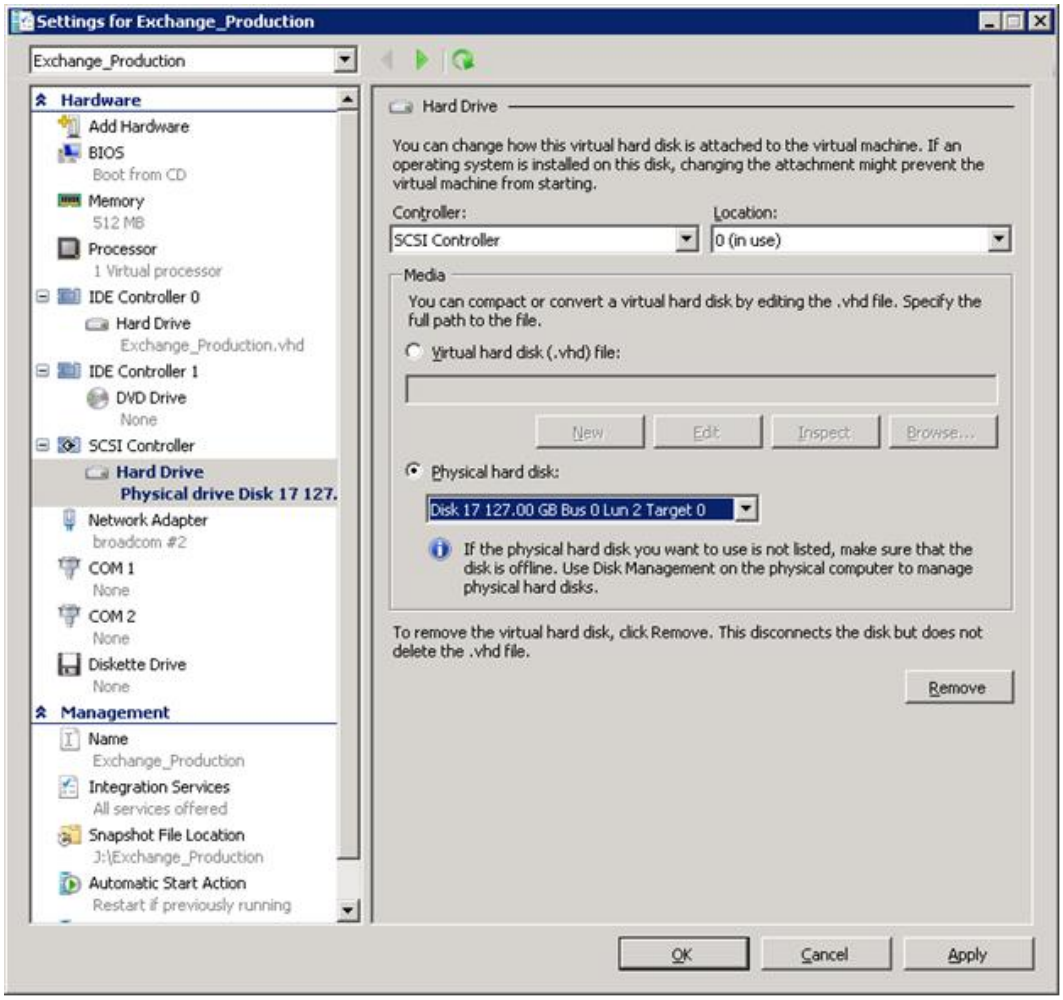

This presents a new drive inside the VM (backed by VHD in the parent). The user can then log in to the VM and choose the newly mounted drive and see the contents of the disk backed by the VHD attached.

After the verification is done, detach the disk from the virtual machine using the VM settings page and choose the Remove button. Use SnapDrive for Windows to unmount the disk backed by the Snapshot copy, using the SnapDrive Disconnect disk MMC action/menu item. Customers can also use SDCLI Snap Unmount command to unmount the disk mounted from Snapshot technology.

Note: Leaving a backup in a mounted state places Snapshot copies in a busy condition, preventing the deletion of both the mounted backup and any preceding Snapshot copies. Backup should be unmounted when not in use.

7 SnapManager for Hyper-V High Availability

The availability of the shared storage infrastructure is more critical than the actual availability of the individual physical servers hosting the VMs on a Hyper-V server itself because they support features such as live/quick migration, which makes sure of the high availability at the hypervisor layer. With the NetApp software solution, most of the availability requirements of a virtual infrastructure can be addressed.

Note that the SMHV, as a host-end application, offers services provided that the storage is continuously available. Following is a detailed description of the available tools that facilitate storage availability.

7.1 Multipath HA with Active-Active NetApp Controllers

The NetApp active-active controllers offer easy, automatic, and transparent failover capabilities to deliver a high-availability (HA) solution. Configuring multipath HA with NetApp active-active controllers enhances the overall storage infrastructure availability and promotes higher performance consistency. It offers protection against storage failure events such as FC adapter or port failure, controller-to-shelf cable failure, shelf module failure, dual intershelf cable failure, and secondary path failure. This equips environments running business-critical applications such as the Microsoft Hyper-V virtual infrastructure to provide uninterrupted services.

Best Practices
Use active-active storage controller configuration to eliminate any single points of failure (SPOFs). Use multipath HA with active-active storage configuration to get a better storage availability and higher performance. More details on high-availability system configuration can be obtained from NetApp TR-3450: Active-Active Controller Overview and Best Practices Guidelines.

8 SnapManager for Hyper-V Disaster Recovery for SAN Environments

The user can perform failover and failback of Hyper-V VMs using Windows PowerShell cmdlets in the SMHV PowerShell option. The Windows PowerShell cmdlet `restore-backup` must be used along with the switch `-RestoreToAlternateHost` and the server name to use this feature.

For example:

```
PS C:\Windows\system32> restore-backup -server cluster_1 -RestoreToAlternateHost -
disableverifysnapshot -backup DR_Dataset_Secondary_01-22-2010_18.21.33 -resourcename smhv-demo-
csv -verbose
```

Get-VMsFromBackup Cmdlet

This cmdlet is used to retrieve the VMs from backup metadata. In a DR scenario, the administrator has access to the backup metadata from the primary site and must know which VMs are present in the backup in order to restore them on the secondary site. This new cmdlet provides a list of VMs present in the backup.

The `-server` switch of this cmdlet is used to specify the hostname or cluster name on the secondary site. SMHV looks for the backups in SnapInfo for this input host or cluster and finds the VMs present in these backups.

For example:

```
PS C:\Windows\system32> get-vmsfrombackup -server cluster_windows2008_r2
Name Id
SMHV-demo-CSV F10F1011-901A-4789-ADE4-A1F34323E2D7
```

Prerequisites

- Site A (primary) containing storage systems and standalone Hyper-V host system or Hyper-V host cluster. VMs running on these hosts reside on NetApp storage.
- Site B (secondary) containing storage systems and Hyper-V host or cluster (same as that of primary site).
- SDW and SMHV are installed on both site A and site B.
- A SnapMirror relationship is initialized from site A to site B.

- A Hyper-V host or cluster on site A is added to SMHV, and the VMs are backed up using SMHV. The policy to update SnapMirror after backup is checked. Thus, after each backup, the secondary site is updated with new Snapshot copies of VMs and SnapInfo.

Steps to Fail Over VMs to Secondary Site

Following are the steps to fail over VMs to secondary site:

1. Connect to all of the LUNs from secondary storage system volumes. If the secondary site is a cluster, go to the node where a cluster group is online and connect to all of the LUNs from that node in the cluster. SDW breaks the SnapMirror relationship and also uses SnapRestore. If the volume contains only one LUN, SDW performs a volume-based SnapRestore (VBSR) operation, and the SnapMirror relationship is then in an uninitialized state. If the volume contains multiple LUNs, SDW performs a single-file SnapRestore operation, and the SnapMirror relationship is broken off.
2. Restore the SnapInfo LUN from the last Snapshot copy of it that was created by SMHV.
3. Add the secondary host or cluster in SMHV and configure it with the SnapInfo path.
4. Use the Get-VMsFromBackup cmdlet to get a list of the VMs present in backup metadata.
5. Use the Get-Backup cmdlet to get the backups for each VM and the list of files in the Snapshot copy.
6. Use the Restore-backup cmdlet with VM GUID (from step 4), backup (from step 5), and list of VHDs (from step 5). Use the -RestoreToAlternateHost switch and specify the secondary host or cluster name as -server parameter. If the secondary site is a cluster, make sure that the LUNs on which VMs reside are online on the cluster node that owns the cluster group.

If the secondary site is a cluster, make VMs highly available using failover cluster UI/Windows PowerShell cmdlets.

Example

```
Restore-Backup -ResourceId 05391048-68F5-4153-84F6-52C9643F4592 -RestoreToAlternateHost -Verbose
-DisableVerifySnapshot -VirtualMachinePath "K:\testVM" -SnapshotFilePath "K:\testVM" -VHDs
@(@{"SourceFilePath" = "J:\testVM\test1.vhdx"; "DestinationFilePath" = "K:\testVM\test1.vhdx"},
@{"SourceFilePath" = "J:\test.vhdx"; "DestinationFilePath" = "K:\test.vhdx"}, @{"SourceFilePath"
= "J:\testVM\VirtualHardDisks\test2.vhdx"; "DestinationFilePath" =
"K:\testVM\VirtualHardDisks\test2.vhdx"}, @{"SourceFilePath" = "J:\testVM\Virtual Hard
Disks\test3.vhdx"; "DestinationFilePath" = "K:\testVM\Virtual Hard Disks\test3.vhdx"}) -
BackupName sm_san_ded1_07-10-2013_11.49.06
```

In the preceding example, a VM called "testVM" having four VHDs is stored in primary host in J: drive. It is restored to the K: drive in the secondary host. K:\ is mapped to secondary mirrored volume using SnapMirror.

If the secondary site is an active site with its own virtual machine LUNs and SnapInfo LUN, then in order to restore the VMs present in the primary site to the secondary site:

7. Connect the primary SnapInfo LUN to the secondary host by breaking the mirrored volume.
8. Use SnapRestore from the last SMHV SnapInfo Snapshot copy.
9. Copy the contents to the already existing SnapInfo copy to the secondary host.

In this manner, the VMs in the primary site are reflected in the SMHV console of the secondary site and can be managed appropriately.

Steps to Fail Back VMs to Primary Site

Following are the steps to fail back VMs to primary site:

1. Get the data from the secondary site back onto the primary storage system.
2. If the primary site is completely destroyed, new storage must be provisioned. If that is done, the user must initialize the SnapMirror relationship from the secondary storage system to the primary storage system (this is a new relationship) to get the data back. After the relationship is initialized and the data is back on the primary, this relationship can be released.

3. If the primary site was temporarily down, the user must This is confusing. “the user must resync the primary site with only those changes that were implemented on the secondary site while the primary site was gone. To do this, resync existing SnapMirror relationships in the reverse direction (resync from the secondary site to the primary site).
4. When the data on the secondary storage system is synchronized with the primary storage system, go to the SnapDrive user interface (UI) on the secondary storage system and initiate a SnapMirror update for each of the LUNs on the secondary storage system. If this is not done, SDW uses the SMHV backup Snapshot copy to restore the LUNs on the primary storage system. The LUN in the backup Snapshot copy is actually a LUN clone, so this must be avoided by forcing one more SnapMirror updates.

   Note: Taking an SMHV backup (with the Update SnapMirror option checked) from the secondary storage system has the same effect as manually doing the SnapMirror update from the SDW graphical user interface (GUI). Most users will probably create the SMHV backup in lieu of manually performing a mirror update because it can be scripted, whereas the mirror update is a tedious job and prone to user error (such as forgetting to update a LUN).
5. Connect to all LUNs on the primary site (same type, same mount points). If the primary is a cluster, go to the node where the cluster group is online and connect to all the LUNs from that node in the cluster. If a resync in reverse direction has been done, there will be a new broken (or uninitialized) SnapMirror relationship from the secondary storage system to the primary storage system. This can be released.
6. Restore the SnapInfo LUN from its last Snapshot copy created by SMHV.
7. Add the primary host or cluster in SMHV MMC and configure it with the SnapInfo path.
8. Use the Get-VMsFromBackup cmdlet to get a list of VMs present in backup metadata.
9. Use the Get-Backup cmdlet to get the backups for each VM.
10. Use the Restore-backup cmdlet with VM GUID (from step 6) and backup (from step 7). Use -RestoreToAlternateHost switch and specify the primary host or cluster name as -server parameter. If the primary host is a cluster, make sure that the LUNs (cluster resources) on which the VM resides are online on the node that owns the cluster group.
11. If the primary host is the cluster, make the VMs highly available using failover cluster UI/Windows PowerShell cmdlets.

After the VMs are backed up on the primary site, it is necessary to get back to the original configuration with a SnapMirror relationship established from the primary storage system to the secondary storage system. To do this, perform the following steps on the secondary site:

1. If the secondary site is a standalone host, shut down and delete the VMs running on it. Disconnect the SnapInfo disk and the disks containing VMs using SnapDrive. If the secondary host is a cluster, offline the VM resource and VM configuration resource for all the VMs. Delete these resources from the cluster. Delete all the VMs from Hyper-V Manager. Disconnect all disks using SnapDrive.
2. Resync the SnapMirror relationship from the primary storage system to the secondary storage system.

9 Conclusion

SnapManager for Hyper-V provides a rich feature set that allows IT organizations to take advantage of NetApp Snapshot and SnapMirror technologies to provide fast, space-efficient disk-based backups in a Hyper-V environment with NetApp storage while placing minimal overhead on the associated virtual infrastructure. The recommendations and examples in this report will help administrators get the most out of SMHV deployments.

Appendix

How to Choose Your Hyper-V and VHD Storage Container Format

Customers have to make a choice when they need to decide what the appropriate storage container format is for deploying virtual machines using Hyper-V.

The summary in Table 3 is intended to make the decision-making process easier.

Table 3) Choosing Hyper-V and VHD storage container format.

Storage Container	Pros	Cons
Pass-through disk	• Fastest performance • Simplest storage path because file system on host is not involved • Better alignment under SAN • For shared storage based pass-through, no need to mount the file system on host and that might speed up VM live migration • Lower CPU use • Support very large disks	• VM Snapshot copy cannot be created • Disk is being used exclusively and directly by a single virtual machine • Pass-through disks cannot be backed up by the Hyper-V VSS writer and any backup program that uses the Hyper-V VSS writer
Fixed size VHD	• Highest performance of all VHD types • Simplest VHD file format to give the best I/O alignment • More robust than dynamic or differencing VHD because of the lack of block allocation tables (redirection layer) • File-based storage container has more management advantages than pass-through disk • Expanding is available to increase the capacity of VHD • No risk of underlying volume running out of space during VM operations	• Up-front space allocation might increase the storage cost when large number of fixed VHDs is deployed • Large fixed VHD creation is time-consuming • Shrinking the virtual capacity (reducing the virtual size) is not possible
Dynamically expanding or differencing VHD	• Good performance • Quicker to create than fixed-size VHD • Grow dynamically to save disk space and provide efficient storage usage • Smaller VHD file size makes it more nimble in terms of transporting across the network • Blocks of full zeros will not get allocated and thus save the space under certain circumstances • Compact operation is available to reduce the actual physical file size	• Interleaving of metadata and data blocks might cause I/O alignment issues • Write performance might suffer during VHD expanding • Dynamically expanding and differencing VHDs cannot exceed 2040GB • Might get VM paused or pull the VHD out if disk space is running out due to the dynamic growth • Shrinking the virtual capacity is not supported • Expanding is not available for differencing VHDs due to the inherent size limitation of parent disk

Storage Container	Pros	Cons
		• Defrag is not recommended because of inherent redirectional layer

SMHV: Virtual Machine Self-Management

If a VM belongs a host that has SMHV installed, and you install SMHV on that VM to use as a management console, you should not use SMHV to manage the host to which the VM belongs.

For example, if VM1 belongs to Host1 (with SMHV installed), and you install SMHV on VM1, you should not use SMHV to manage Host1 from VM1.

If you do this and try to restore the VM from itself, the VM will be deleted or restarted from Hyper-V Manager.

SMHV: Data ONTAP VSS Hardware Provider Requirement

Data ONTAP VSS hardware provider must be installed for SnapManager to function properly. Data ONTAP VSS hardware provider integrates the SnapDrive service and storage systems running Data ONTAP into the VSS framework. The Data ONTAP VSS hardware provider is now included with SnapDrive 6.0 or later and does not need to be installed separately.

Viewing Installed VSS Providers

1. To view the VSS providers installed on your host, complete these steps:
2. Select Start.
3. Run and enter the following command to open a Windows command prompt: cmd.
4. At the prompt, enter the following command:

```
Vssadminlist providers
```

The output should be similar to the following:

```
Provider name: 'Data ONTAP VSS
Hardware Provider' Provider type:
Hardware
Provider Id: {ddd3d232-a96f-4ac5-8f7b-250fd91fd102}
Version: 6.2.0.xxxx
```

Verifying That VSS Hardware Provider Was Used Successfully

To verify that the Data ONTAP VSS hardware provider was used successfully after a Snapshot copy was created, complete this step.

Navigate to System Tools > Event Viewer > Application in MMC and look for an event with the following values:

```
Source Event ID Description
The VSS provider has successfully completed CommitSnapshots for SnashotSetId id in n
milliseconds. Navsspr 4089
```

Note: VSS requires that the provider initiate a Snapshot copy within 10 seconds. If this time limit is exceeded, the Data ONTAP VSS hardware provider logs event ID 4364. This limit could be exceeded because of a transient problem. If this event is logged for a failed backup, retry the backup.

SMHV: When Virtual Machine Backups Take Too Long to Complete

If a virtual machine contains several direct-attached iSCSI LUNs or pass-through LUNs, and SnapDrive for Windows is installed on the virtual machine, the virtual machine backup can take a long time. The Hyper-V writer takes a hardware Snapshot copy of all the LUNs in the virtual machine using the SnapDrive for Windows VSS hardware provider. There is a Microsoft hotfix that uses the default system provider (software provider) in the virtual machine to make the Snapshot copy. As a result, the Data ONTAP VSS hardware provider is not used for Snapshot copy creation inside the child OS, and the backup speed increases. For more information on the Microsoft hotfix, see Knowledge Base article 975354 on the Microsoft support site at http://support.microsoft.com/.

SMHV: Redirected I/O and VM Design Considerations

Although redirected I/O is handled in a Windows Server 2008R2 Hyper-V cluster, SMB API calls are made from one cluster node to the cluster and CSV owner. This involves metadata traffic and other SMB API calls that can affect performance significantly.

NetApp recommends that the user manually assign CSV and VM ownership to specific nodes in the cluster. SMHV backup datasets must be created and designed to back up all VMs in a single CSV owned by each specific node as follows:

1. Using SnapDrive for Windows, create one CSV per host cluster node, based upon tiers of storage as necessary. For example, create one CSV for fast SAS disk and one for SATA.
2. Using SCVMM, migrate VMs into their respective CSVs and assign ownership of those VMs to the same node that owns the CSV.

Note: All VM migrations should be performed using SCVMM.

3. Create an SMHV dataset for each CSV and make sure that all VMs that reside in that CSV are placed into that dataset. For best results, do not allow VMs owned by multiple nodes to coreside within the same CSV.
4. Create a backup policy for each dataset that matches the customer's backup needs.
5. Using Failover Cluster Manager:
    a. Assign preferred ownership of each VM to its appropriate cluster node.
    b. Assign preferred ownership of each CSV to its appropriate cluster node.
    c. Before running each backup for each cluster node, assign cluster master ownership to the cluster node being backed up by that SMHV dataset. This is done through Failover Cluster Manager or using a Windows PowerShell script that can be executed by SMHV at the beginning of the backup job.

SMHV: Transferring Snapshot Copies to SnapVault or a Tape Device

In order to transfer SMHV Snapshot copies to SnapVault or a tape device, users can create a script and use the SMHV postscript feature in SMHV dataset policy option.

SMHV offers the following predefined variables, which the administrator can pass in order to achieve this:

- $VMSnapshot
- $SnapInfoName
- $SnapInfoSnapshot

During the post–policy execution phase, SMHV will replace the $VMSnapshot variable with the Snapshot name, $SnapInfoName with the time stamp of the backup, and $SnapInfoSnapshot with the snapinfo Snapshot name. You can access these variables from your scripts and do the necessary actions.

Here is a sample script that transfers SMHV Snapshot copies to a SnapVault system:

The following scripts are used:

- sv_update.ps1: Is used to update the Hyper-V VM Snapshot copies to secondary storage
- sv_update_snapinfo.ps1: Is used to update SnapManager for Hyper-V to secondary storage
- update-vmsnapshot.bat: Batch file is used to start the sv_update.ps1 file with the correct parameters
- update-snapinfo.bat: Batch file is used to start the sv_update_snapinfo.ps1 file with the correct parameters

Prerequisites

1. Download and install the NetApp Data ONTAP PowerShell Toolkit: https://communities.netapp.com/community/products_and_solutions/microsoft/powershell/data_ontap_powershell_toolkit_downloads.
2. Unzip the PowerShell Toolkit to C:\Windows\System32\WindowsPowerShell\v1.0\Modules.
3. In Windows PowerShell, set the Set-ExecutionPolicy to "RemoteSigned"; otherwise, no scripts are allowed to be run. This needs to be done on every Hyper-V host.

Configuration Procedure for SnapVault Script

- Primary – test02
- Secondary – test03

1. Create SnapVault relationships:

```
snapvault start -S test02:/vol/csv01/csv01 test03:/vol/csv01/csv01
snapvault start -S test02:/vol/csv02/csv02 test03:/vol/csv02/csv02
snapvault start -S test02:/vol/csv03/csv03 test03:/vol/csv03/csv03
snapvault start -S test02:/vol/smhv_snapinfo/snapinfo test03:/vol/smhv_snapinfo/snapinfo
```

2. Set retention time on the secondary volumes:

   From test03:

```
snapvault snap sched csv01 sv_daily_testhv01 7@-
snapvault snap sched csv01 sv_daily_testhv02 7@-
snapvault snap sched csv01 sv_daily_testhv03 7@-
snapvault snap sched csv01 sv_daily_testhv04 7@-
snapvault snap sched csv01 sv_daily_testhv05 7@-
snapvault snap sched csv01 sv_daily_testhv06 7@-
snapvault snap sched csv01 sv_daily_testhv07 7@-

snapvault snap sched csv02 sv_daily_testhv01 7@-
snapvault snap sched csv02 sv_daily_testhv02 7@-
snapvault snap sched csv02 sv_daily_testhv03 7@-
snapvault snap sched csv02 sv_daily_testhv04 7@-
snapvault snap sched csv02 sv_daily_testhv05 7@-
snapvault snap sched csv02 sv_daily_testhv06 7@-
snapvault snap sched csv02 sv_daily_testhv07 7@-

snapvault snap sched csv03 sv_daily_testhv01 7@-
snapvault snap sched csv03 sv_daily_testhv02 7@-
snapvault snap sched csv03 sv_daily_testhv03 7@-
snapvault snap sched csv03 sv_daily_testhv04 7@-
snapvault snap sched csv03 sv_daily_testhv05 7@-
snapvault snap sched csv03 sv_daily_testhv06 7@-
snapvault snap sched csv03 sv_daily_testhv07 7@-

snapvault snap sched smhv_snapinfo  sv_daily 7@-
```

3. Enable SIS on the secondary volumes and start first dedupe run:

```
sis on /vol/csv01
sis on /vol/csv02
sis on /vol/csv03
sis on /vol/smhv_snapinfo
```

```
sis start -s /vol/csv01
sis start -s /vol/csv02
sis start -s /vol/csv03
sis start -s /vol/smhv_snapinfo
```

4. Create a "scripts" folder under C:\Program Files\NetApp\SnapManager for Hyper-V.

   This needs to be done on every Hyper-V host.

5. Place the Windows PowerShell scripts and batch files into the "C:\Program Files\NetApp\SnapManager for Hyper-V\scripts\" folder of every SMHV server.

   This needs to be done on every Hyper-V host.

6. Configure SMHV.
7. Create an SMHV dataset.
8. Create a policy for the dataset.
   a. Add the postscript "update-vmsnapshot.bat" (batch files that call the Windows PowerShell script) to the policy.
   b. Also, add the parameters "`$VMSnapshot $SnapInfoName`" to the "Arguments" box.
9. Try the backup and review the result.
   a. Now, all the VM LUNs will have a backup with a consistent copy to SnapVault. Next step is to create a schedule for the snapinfo LUN.
10. Create a Windows task, which will kick off the snapinfo update script. Schedule this to run after the SMHV backup. Make sure it runs when the SMHV backup is finished. Otherwise, there will be no snapinfo Snapshot copy. The snapinfo Snapshot copy is only created after all the nodes of the SMHV cluster are finished with the backup.

Script 1: sv_update.ps1

```
if($ARGS.Length -lt 8)
{
   cls
   write $(" ");
   write $("Usage: sv_update.ps1 <primary_filer> <secondary_filer> <primary_volume>
<secondary_volume> <secondary_path> <retention_period> ");
   write $(" ");
   write $("Example: sv_update.ps1 filer01 filer02 vol1 sv_vol1 /vol/sv_vol1/qtree1 sv_daily ");
   write $(" ");
   exit(1);
}

$prifiler = $ARGS[0]
$secfiler = $ARGS[1]
$privol = $ARGS[2]
$secvol = $ARGS[3]
$secpath = $ARGS[4]
$sret = $ARGS[5]
$VMSnapshot = $ARGS[6]
$SnapInfoName = $ARGS[7]

#Reformat the VMSnapshot string and SnapINfo Snapshot string
$VMSnapshot_backup = $VMSnapshot+"_backup"

#Get the machine name of the HV host in lowercase for the retention period
$hvhost = $env:computername.ToLower()

#set retention period for this HV Host
$hvhostret = $sret+"_"+$hvhost

Import-Module dataontap


#check to see if the snapshot exist on the primary volume. If not we will exit the script.
Connect-NaController $prifiler
```

```
$str = Get-NaSnapshot -Targetname $privol -snapname $VMSnapshot
if ($str)
{
        # 1 - Initiates SnapVault transfer (update) from Secondary using last SnapManager for
Hyper-v snapshot.
        Connect-NaController $secfiler
        Start-NaSnapVaultSecTransfer $secpath -PrimarySnapshot $VMSnapshot

        # 2 - Simple time loop that will wait until SV update on Secondary (snapvault_secondary
volume) is done, before creating snapshot (Step 4).
        #     This script loops every 5 seconds until SnapVault status shows "Idle". Steps 5 & 6
should run once a month against Secondary.
        Connect-NaController $secfiler
        $var = $null
        while (!$var -or ($var.status -ne "Idle"))
        {
          $var = Get-NaSnapvaultSecStatus -Path $secpath
          start-sleep -seconds 5
        }

        # 4 - Initiates SnapVault transfer (update) from Secondary using last SnapManager for
Hyper-v snapshot.
        # This is the application persistent snapshot

        Connect-NaController $secfiler
        Start-NaSnapVaultSecTransfer $secpath -PrimarySnapshot $VMSnapshot_backup -
NoLunCloneExpansion 1

        # 5 - Simple time loop that will wait until SV update on Secondary (snapvault_secondary
volume) is done, before creating snapshot (Step 4).
        #     This script loops every 5 seconds until SnapVault status shows "Idle". Steps 5 & 6
should run once a month against Secondary.

        Connect-NaController $secfiler
        $var = $null
        while (!$var -or ($var.status -ne "Idle"))
        {
          $var = Get-NaSnapvaultSecStatus -Path $secpath
          start-sleep -seconds 5
        }

        # 4 - This archives (creates snapshot) on Secondary using the given retention schedule.
        Connect-NaController $secfiler
        Start-NaSnapvaultSecSnapshot -VolumeName $secvol -ScheduleName $hvhostret
exit(1);
}

else
{
  write $("Nothing to do there is no primary snapshot ");
exit(1)}
```

Script 2: sv_update_snapinfo.ps1

```
if($ARGS.Length -lt 6)
{
   cls
   write $(" ");
   write $("Usage: sv_update.ps1 <primary_filer> <secondary_filer> <primary_volume>
<secondary_volume> <secondary_path> <retention_period>");
   write $("Example: sv_update.ps1 filer01 filer02 vol1 sv_vol1 /vol/sv_vol1/qtree1 sv_daily");
   write $(" ");
   exit(1);
}

$prifiler = $ARGS[0]
$secfiler = $ARGS[1]
$privol = $ARGS[2]
$secvol = $ARGS[3]
$secpath = $ARGS[4]
```

```
$sret = $ARGS[5]


Import-Module dataontap
#check to see if the snapshot exist on the primary volume. If not we will exit the script.
Connect-NaController $prifiler
$str = Get-NaSnapshot -Targetname $privol -snapname $VMSnapshot
if ($str)
{
        # 1 - Pulls the last SnapManager for Hyper-v snapshot (using "smhv_snapinfo" key word)
from the snapvault_primary volume on the Primary.
        Connect-NaController $prifiler
        $LastSnapshot = get-nasnapshot $privol | ? { $_.Name -match "smhv_snapinfo" } | Sort-
Object AccessTimeDT -Descending | Select-Object -first 1

        # 2 - Initiates SnapVault transfer (update) from Secondary using last SnapManager for
Hyper-v snapshot.
        Connect-NaController $secfiler
        Start-NaSnapVaultSecTransfer $secpath -PrimarySnapshot $LastSnapshot.Name

        # 3 - Simple time loop that will wait until SV update on Secondary (snapvault_secondary
volume) is done, before creating snapshot (Step 4).
        #     This script loops every 5 seconds until SnapVault status shows "Idle". Steps 5 & 6
should run once a month against Secondary.
        Connect-NaController $secfiler
        $var = $null
        while (!$var -or ($var.status -ne "Idle"))
        {
          $var = Get-NaSnapvaultSecStatus -Path $secpath
          start-sleep -seconds 5
        }
        # 4 - This archives (creates snapshot) on Secondary using the given retention schedule.
        Connect-NaController $secfiler
        Start-NaSnapvaultSecSnapshot -VolumeName $secvol -ScheduleName $sret
exit(1);
}
else
{
  write $("Nothing to do there is no primary snapshot ");
exit(1)}
```

Script 3: update-vmsnapshot.bat (Batch File)

```
powershell -file "C:\Program Files\NetApp\SnapManager for Hyper-V\scripts"\sv_update_snapinfo.ps1
test02 test03 smhv_snapinfo smhv_snapinfo /vol/smhv_snapinfo/snapinfo sv_daily
```

Script 4: update-snapinfo.bat (Batch File)

```
powershell -file "C:\Program Files\NetApp\SnapManager for Hyper-V\scripts"\sv_update.ps1 test02
test03 csv01 csv01 /vol/csv01/csv01 sv_daily backup %1 %2
powershell -file "C:\Program Files\NetApp\SnapManager for Hyper-V\scripts"\sv_update.ps1 test02
test03 csv02 csv02 /vol/csv02/csv02 sv_daily backup %1 %2
powershell -file "C:\Program Files\NetApp\SnapManager for Hyper-V\scripts"\sv_update.ps1 test02
test03 csv03 csv03 /vol/csv03/csv03 sv_daily backup %1 %2
```

If the user does not want to use a postscript to send the Snapshot copy to the SnapVault or tape system, the user can use this script separately at a different time. In order to do this, we need to sort the Snapshot copies and select the latest Snapshot copy that needs to be sent to the secondary storage system. This can be done invoking the '`snap list`' command in Data ONTAP.

References

- Virtualization with Hyper-V: Supported Guest Operating Systems
  www.microsoft.com/windowsserver2008/en/us/hyperv-supported-guest-os.aspx
- Install the Hyper-V Role on a Full Installation of Windows Server 2008
  http://technet.microsoft.com/en-us/library/cc794929(WS.10).aspx
- NetApp TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide
  www.netapp.com/us/library/technical-reports/TR-3701.html
- Install the Hyper-V Role on a Server Core Installation of Windows Server 2008
  http://technet.microsoft.com/en-us/library/cc794852(WS.10).aspx
- Microsoft Hyper-V Server 2008 Configuration Guide
  www.microsoft.com/Downloads/details.aspx?familyid=E1E111C9-FA69-4B4D-8963-1DD87804C04F&displaylang=en
- Infrastructure Planning and Design Guide for System Center Virtual Machine Manager 2008 R2
  http://technet.microsoft.com/en-us/library/cc196387.aspx
- Nvspbind
  http://archive.msdn.microsoft.com/nvspbind
- VMM System Requirements
  http://technet.microsoft.com/en-us/library/cc764328.aspx
- Configuring a SAN Environment for VMM
  http://technet.microsoft.com/en-us/library/cc764269.aspx
- New Installation of VMM
  http://technet.microsoft.com/en-us/library/cc793149.aspx
- Configuring Virtual Networks
  http://technet.microsoft.com/en-us/library/cc816585(WS.10).aspx
- Hyper-V: Using Hyper-V and Failover Clustering
  http://technet.microsoft.com/en-us/library/cc732181(WS.10).aspx
- Hyper-V: Using Live Migration with Cluster Shared Volumes in Windows Server 2008 R2
  http://technet.microsoft.com/en-us/library/dd446679(WS.10).aspx
- Hyper-V Live Migration Overview and Architecture
  www.microsoft.com/downloads/details.aspx?FamilyID=FDD083C6-3FC7-470B-8569-7E6A19FB0FDF&displaylang=en
- Virtual Network Manager
  http://technet.microsoft.com/en-us/library/cc754263.aspx
- New in Hyper-V Windows Server 2008 R2 Part 1: Dedicated Networks
  http://blogs.technet.com/jhoward/archive/2009/05/04/new-in-hyper-v-windows-server-2008-r2- part-1-dedicated-networks.aspx
- NetApp Support
  http://support.netapp.com/
- NetApp Interoperability Matrix
  https://now.netapp.com/matrix/mtx/login.do
- High-Availability System Configuration
  www.netapp.com/us/products/platform-os/active-active.html
- NetApp TR-3450: Active-Active Controller Overview and Best Practices Guidelines
  www.netapp.com/us/library/technical-reports/tr-3450.html
- NetApp TR-3437: Storage Subsystem Resiliency Guide
  www.netapp.com/us/library/technical-reports/tr-3437.html
- NetApp Fibre Channel and iSCSI Configuration Guide
  http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/

- Remote LAN Management (RLM)
  http://now.netapp.com/NOW/download/tools/rlm_fw/info.shtml
- Windows Host Utilities 5.0 Installation and Setup Guide
  http://now.netapp.com/NOW/knowledge/docs/san/#windows
- Data ONTAP DSM for Windows MPIO Installation and Administration Guide
  http://now.netapp.com/NOW/knowledge/docs/san/#mpio
- Data ONTAP Network and File Access Management Guide
  http://now.netapp.com/NOW/knowledge/docs/ontap/
- NetApp SnapDrive for Windows
  www.netapp.com/us/products/management-software/snapdrive-windows.html
- NetApp SnapManager Family
  www.netapp.com/us/products/management-software/
- SnapManager for Hyper-V Installation and Administration Guide
  http://support.netapp.com/documentation/productlibrary/index.html?productID=30055
- Performance Tuning Guidelines for Windows Server 2008 R2
  www.microsoft.com/whdc/system/sysperf/Perf_tun_srv-R2.mspx
- What's New in Windows Server 2008 R2 Hyper-V Performance and Scale?
  http://blogs.msdn.com/tvoellm/archive/2009/08/05/what-s-new-in-windows-server-2008-r2- hyper-v-performance-and-scale.aspx
- Microsoft Virtual Hard Disk (VHD) FAQ
  http://technet.microsoft.com/en-us/bb738381.aspx
- Virtual Hard Disk (VHD) Image Format Specification
  http://technet.microsoft.com/en-us/virtualserver/bb676673.aspx
- Performance Tuning Guidelines for Windows Server 2008 R2
  www.microsoft.com/whdc/system/sysperf/Perf_tun_srv.mspx
- Hyper-V and VHD Performance: Dynamic vs. Fixed
  http://blogs.technet.com/winserverperformance/archive/2008/09/19/hyper-v-and-vhd- performance-dynamic-vs-fixed.aspx
- Data ONTAP Block Access Management Guide for iSCSI or FC
  http://now.netapp.com/NOW/knowledge/docs/ontap/rel731_vs/pdfs/ontap/bsag.pdf
- Data ONTAP Commands Manual Page Reference, Volumes 1 and 2
  https://now.netapp.com/AskNOW/search?action=search&search_collections=kbase&search_collections=bugs&search_collections=docs&search_collections=tools&query=Data%20ONTAP%20Commands%20Manual%20Page%20Reference
- Planning for Disks and Storage
  http://technet.microsoft.com/en-us/library/dd183729(WS.10).aspx
- Configuring Disks and Storage
  http://technet.microsoft.com/en-us/library/ee344823(WS.10).aspx
- NetApp TR-3505: NetApp Deduplication for FAS: Deployment and Implementation Guide
  www.netapp.com/us/library/technical-reports/tr-3505.html
- Microsoft KB 302577: http://support.microsoft.com/kb/302577
- Microsoft KB 958184: http://support.microsoft.com/kb/958184
- Microsoft KB 961804: http://support.microsoft.com/kb/961804/en-us
- SnapMirror Sync and SnapMirror Semi-Sync Overview and Design Considerations
  www.netapp.com/us/library/technical-reports/tr-3326.html
- SnapMirror Async Overview and Best Practices Guide
  www.netapp.com/us/library/technical-reports/tr-3446.html

- Configuring Alarms
  http://now.netapp.com/NOW/knowledge/docs/DFM_win/rel36r1/html/software/opsmgr/monitor5.htm
- Managing Aggregate Capacity
  http://now.netapp.com/NOW/knowledge/docs/DFM_win/rel36r1/html/software/opsmgr/filesys4.htm
- Operations Manager
  www.netapp.com/us/products/management-software/operations-manager.html
- A Description of the Diskpart Command-Line Utility
  http://support.microsoft.com/default.aspx?scid=kb;en-us;300415
- GNU ext2resize
  http://ext2resize.sourceforge.net/
- NetApp Data ONTAP PowerShell Toolkit
  http://communities.netapp.com/docs/DOC-6162#
- Data ONTAP Block Access Management Guide
  http://now.netapp.com/NOW/knowledge/docs/ontap/rel733/pdfs/ontap/bsag.pdf
- SnapDrive for Windows Installation and Administration Guide
  http://now.netapp.com/NOW/knowledge/docs/snapdrive/relsnap63/pdfs/admin.pdf
- Data ONTAP 7.3 System Administration Guide
  http://now.netapp.com/NOW/knowledge/docs/ontap/rel733/pdfs/ontap/sysadmin.pdf

Knowledge Base Articles

- KB ID: 3011206: SMHV: Can SnapManager 1.0 for Hyper-V exclude Virtual Hard Disks from backups?
- KB ID: 1010146: SMHV: How to manually restore a Hyper-V virtual machine from a Snapshot backup
- KB ID: 1011587: How to migrate a Hyper-V VM to support SnapManager for Hyper-V Backup
- KB ID: 2010899: SMHV: Backups fail for Hyper-V Virtual Machines containing Passthru or iSCSI in Guest Disks
- KB ID: 1010887: SMHV: How to setup SnapInfo Logical Unit Number (LUN)
- KB ID: 2010607: SMHV: Creation of two snapshots for every backup
- KB ID: 2014905: SnapManager for Hyper-V backups fail to complete even though all Virtual Machines are located on NetApp LUNs
- KB ID: 2014900: SnapManager for Hyper-V backup sets that contain Windows XP fail
- KB ID: 2014928: SMHV: During backup of CSV, hosts report `NO_DIRECT_IO_DUE_TO_FAILURE`.
- KB ID: 2014933: SMHV: Cluster Shared Volume goes offline after backup
- KB ID: 2639032: "0x0000003B," "0x00000027," and "0x0000007e" Stop errors when a connection to a CSV is lost on a Windows Server 2008 R2-based failover cluster
- KB2517329: Performance decreases in Windows Server 2008 R2 when the Hyper-V role is installed on a computer that uses Intel Westmere or Sandy Bridge processors
- KB2552040: A Windows Server 2008 R2 failover cluster loses quorum when an asymmetric communication failure occurs
- KB2522766: The MPIO driver fails over all paths incorrectly when a transient single failure occurs in Windows Server 2008 or in Windows Server 2008 R2
- KB2528357: Nonpaged pool leak when you disable and enable some storage controllers in Windows 7 or in Windows Server 2008 R2
- KB2770917: This is a Windows Server 2012 KB fix to fix the following error:
  “Error: Vss Requestor - Backup Components failed. Writer Microsoft Hyper-V VSS Writer involved in backup or restore encountered a retryable error. Writer returned failure code 0x800423f3. Writer state is 8.” This issue is caused by inclusion of direct-attached iSCSI LUNs or pass-through disks in the VSS backups.

Version History

Version	Date	Document Revision History
1.0	October 2013	Initial release

Acknowledgements

The author would like to thank the following people for their contributions:

- Anagha Barve, member of technical staff, Microsoft Business Unit
- Atul Bhalodia, senior engineer, Microsoft Business Unit
- Chance Bingen, escalation engineer
- Chris A. Collins, enterprise infrastructure architect
- John Fullbright, reference architect, Microsoft Business Unit
- Vineeth Karinta, member of technical staff, Microsoft Business Unit

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®