Technical Report

NetApp SnapManager 2.0 for Hyper-V on Clustered Data ONTAP 8.2

Best Practices Guide

Santhosh Harihara Rao, NetApp
September 2013 | TR-4226

Abstract

This technical report provides guidelines and best practices for integrated architecture 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 environmentally friendly.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Executive Summary

Server virtualization is a major component of data center virtualization and plays a key role in the virtualization initiative. Microsoft is a lead player in this initiative with its server virtualization solutions. This technical report provides detailed guidance on how to architect and implement Microsoft server virtualization solutions on NetApp storage using the clustered Data ONTAP® 8.2 architecture. It describes the use of and best practices for using SnapManager® for Hyper-V (SMHV), a NetApp tool that uses the NetApp Snapshot™ technology for backup, recovery and replication of virtual machines (VMs) in a Hyper-V environment.

NetApp has been on the forefront of solving complex business problems with its innovative technology breakthroughs and end-to-end solutions approach. This technical report is not intended to be a definitive implementation or solutions guide. Expertise might be required to solve specific deployments. Contact your local NetApp sales representative to speak with one of our Microsoft Hyper-V solutions experts. NetApp is dedicated to helping you transform your data center to help your business go further, faster.

2 Scope

This document provides prescriptive guidance and best practices for deploying SnapManager for Hyper-V in a Windows virtual environment. Unless otherwise specified, the best practices described in this document apply to SnapManager 2.0 for Hyper-V installed with SnapDrive® 7.0 for Windows (SDW) in a clustered Data ONTAP 8.2 environment.

3 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, reducing not only the number of physical systems, but also network, power, and administrative overhead.

In contrast to physical environments, in which server resources are underutilized, virtual environments have fewer resources available. Whereas in the past each physical server had dedicated network and CPU resources, VMs must now share those same resources. This can create performance issues, especially while the virtual environment is being backed up, because many VMs use host network and CPU resources concurrently. As a result, backups that once completed during nonbusiness hours are no longer able to finish in their backup window.

NetApp SnapManager for Hyper-V addresses the resource utilization issue typically found in virtual environments by leveraging the underlying NetApp Snapshot technology. This reduces the CPU and network load on the host platforms and drastically reduces 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, allowing quick and efficient restorations, and reducing administrative overhead.

Backups, restores, and disaster recovery (DR) can place a huge overhead on the Hyper-V virtual infrastructure. NetApp SMHV simplifies and automates the backup process by leveraging the underlying NetApp Snapshot and SnapRestore® technologies to provide fast, space-efficient, disk-based backups and rapid, granular restore and recovery of VMs and the associated datasets. The following sections detail the best practices for deploying and using SnapManager for Hyper-V.

3.1 Technical Details

SMHV offers the following capabilities:

- Allows system administrators to create hardware-assisted backup and restore of Hyper-V VMs running on NetApp storage.

- Provides integration with Microsoft Hyper-V VSS writer to quiesce the Hyper-V VMs before creating an application-consistent Snapshot copy of the VM in SAN LUNs.
- Supports backup and restore of VMs running on continuously available SMB shares that are hosted on Data ONTAP 8.2 based systems. Backup operations are performed using a Remote VSS plug-in located in Data ONTAP.
- Allows administrators to create application-consistent backups of Hyper-V VMs if Microsoft Exchange, Microsoft SQL Server®, or any other VSS-aware application is running on VHDs in the VM.
- Provides replication (NetApp SnapMirror®) and vaulting (NetApp SnapVault®) of backup sets to secondary locations for DR planning.
- Supports the backup and restore of shared VMs configured using Windows Failover Clustering (WFC) for high availability (HA) and also on Microsoft Cluster Shared Volumes (CSVs); SMHV supports the seamless processing of scheduled VM backups, regardless of any VM failovers.
- Supports management of multiple remote Hyper-V parent systems from one console.
- Supports performing fast crash-consistent backup and restore of virtual machines.

3.2 Other Reference Documentation

Microsoft Windows Server® 2008 R2 and Windows Server 2012 with the Hyper-V role enabled and coupled with NetApp Data ONTAP offers various storage infrastructure configurations and provisioning methods. In addition to this document, NetApp recommends reading the following documentation before proceeding with SnapManager for Hyper-V deployment:

- TR-3702: NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V
- TR-4172: Microsoft Hyper-V over SMB 3.0 with Clustered Data ONTAP: Best Practices
- 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
- Data ONTAP 8.2 Data Protection Guide

4 SnapManager for Hyper-V Planning

4.1 Storage Considerations

SMHV supports backup and restore of virtual machines on dedicated disks, cluster shared volumes (CSVs), or 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 LUNs for VMs running on shared volumes or CSVs. It does support LUNs created on thin-provisioned volumes and can perform backups and restores on these volumes.

Note: To host VMs in SMB 3.0 shares in Windows Server 2012, the storage system should be running Data ONTAP 8.2

Note: SnapDrive 7.0 for Windows (SDW) must be installed on the host system. NetApp recommends using SnapDrive to provision LUNs or shares to host virtual machines.

Best Practice

For provisioning virtual machines in SMB 3.0 environments, NetApp recommends using the Hyper-V over SMB provisioning templates that are part of the SnapDrive 7.0 for Windows installation. The provisioning templates can be customized based on the environment.

4.2 Product Summary: Supported Platforms and Guidelines

Table 1) SMHV product evolution summary.

SMHV Version	New Features	Supported SnapDrive Version	Supported Data ONTAP Versions
SMHV 1.0	Windows Server 2008 support	SDW 6.3	• Data ONTAP 7-Mode
SMHV 1.0 P1	DR cmdlets	SDW 6.3	• Data ONTAP 7- Mode
SMHV 1.1	Crash-consistent backup	SDW 6.4	• Data ONTAP 7- Mode • Clustered Data ONTAP 8.1
SMHV 1.2	Windows Server 2012 support (CSV 2.0 support)	SDW 6.5	• Data ONTAP 7- Mode • Clustered Data ONTAP 8.1.x
SMHV 2.0	Hyper-V over SMB in Windows Server 2012	SDW 7.0	• Data ONTAP 8.2 (7-Mode and clustered Data ONTAP) • Data ONTAP 8.1.X (7-Mode and clustered Data ONTAP)

Note: Refer to the IMT for full support matrix.

SMHV Deployment Guidelines

- For a version of SMHV, you must install the corresponding supported SnapDrive version indicated in Table 1.
- After upgrading to SMHV 2.0, it is not possible to revert back to the previous version of SMHV.
- To use SMHV in Hyper-V over SMB environments, you must have VMs hosted in Data ONTAP 8.2 systems.
- Before upgrading to the latest SMHV, you must first upgrade to the corresponding SnapDrive version.
- Before upgrading to the latest SnapDrive version, you must upgrade the dependent components, such as Data ONTAP DSM and the Windows Host Utilities Kit. For information about the supported versions of DSM and the Windows Host Utilities Kit, refer to the SnapDrive 7.0 for Windows documentation or the NetApp Interoperability Matrix Tool.
- SDW is required on Hyper-V parent hosts but not on client hosts. For Windows Failover Cluster configurations, SDW and SMHV must be installed on each node of the cluster.
- For backup and restore of virtual machines in SMB 3.0 shares (Hyper-V over SMB 3.0), SnapDrive 7.0 for Windows and SnapManager 2.0 for Hyper-V are required.

License Requirements

An SMHV license is required on the Windows host system. You can choose either host-based or storage-based licensing.

- **Host-based licensing** requires that you enter a license key during installation. To change the license key after installation, click License Settings in the SMHV Welcome window.
- **Storage-based licensing** requires the SMHV license to be added to all storage systems.

The following licenses are required for systems with clustered Data ONTAP 8.2:

- SnapRestore license
- FCP or iSCSI license (for SAN environments)

- CIFS license (for Hyper-V over SMB)
- FlexClone® license (for Hyper-V over SMB)

SnapMirror and SnapVault license are optional.

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, and Enterprise Editions (full and core installation)
- Hyper-V Server 2008 R2 SP1 x64

Remote Management Platform Support

- Windows Server 2012 x64 Standard and Datacenter Editions (full installation)
- Hyper-V Server 2012 R2 x64 (full and core installation)
- Windows Server 2008 x64 Standard and Enterprise Editions (full installation)
- Windows Server 2008 x64 Standard and Enterprise Editions with SP2 (full installation)
- Windows Server 2008 R2 x64 Standard and Enterprise Editions (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

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 installation)
- 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 current information, refer to the NetApp Interoperability Matrix Tool.

Recommended Hotfixes

Table 2) Microsoft hotfixes/updates.

Operating System	Hotfix
Windows Server 2012	2770917
	2779768
	2795944
	2811660
	2822241

Operating System	Hotfix
	2836988
	2845533
	2851998
	2870270
Windows Server 2008 R2	978157
	979743
	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.

5 SnapManager for Hyper-V Architecture

Table 3 lists the terminologies used throughout this document.

Table 3) Terminology.

Term	Description
Dataset	A dataset is a group of VMs that helps to protect data by using retention, scheduling, and replication policies. Datasets can be used to group VMs that have the same protection requirements. A VM can be a member of multiple datasets. This can be useful for VMs that belong to multiple groups (for example, a VM running the SQL Server instance for a Microsoft Office SharePoint® Server [MOSS] configuration might have to belong to both the SQL Server and the MOSS datasets). A dataset cannot have a mix of VMs hosted on both SMB shares and SAN LUNs.
Protection policies	Policies make it possible to schedule or automate the backups of the datasets at a predefined time (schedule policy to provide retention capabilities for older backups [retention policy], and replicate the block changes to the SnapMirror destination volume after the VM backup is created [replication policy]) and also vault the Snapshot copies to a SnapVault destination. Policies include other capabilities that make it possible to run scripts before and after the backup.
Backup and	SMHV provides local backup and recovery capability with the option of replicating backups to a remote storage system by using SnapMirror relationships or updating

Term	Description
recovery	the Snapshot copies to a SnapVault destination. Backups are performed on the whole dataset, which is a logical collection of VMs, with the option of updating the SnapMirror relationship and SnapVault destination as part of the backup on a per-job basis. Similarly, restores can be performed at an individual VM level.
Backup retention policy	Retention policies can be used to specify how long to keep a dataset backup, based on either time or the number of backups. Policies can be created that specify the retention period, allowing administrators the flexibility to meet varying service-level agreements (SLAs) in 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.
Application-consistent backup and restore	These 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 a Snapshot copy. Such a backup guarantees the integrity of application data, and therefore can be safely used to restore the VM and the applications running in the VM to a consistent state.
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 is the same as it would be after a system failure or power outage. This type of backup is much faster than other types. A restore from such a backup is equivalent to a reboot following an abrupt shutdown.

5.1 SMHV Port Usage

For SMHV and SDW, NetApp recommends keeping the following ports open:

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

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

5.2 SMHV 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 (non Hyper-V) servers, and client systems like 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. Once 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 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.

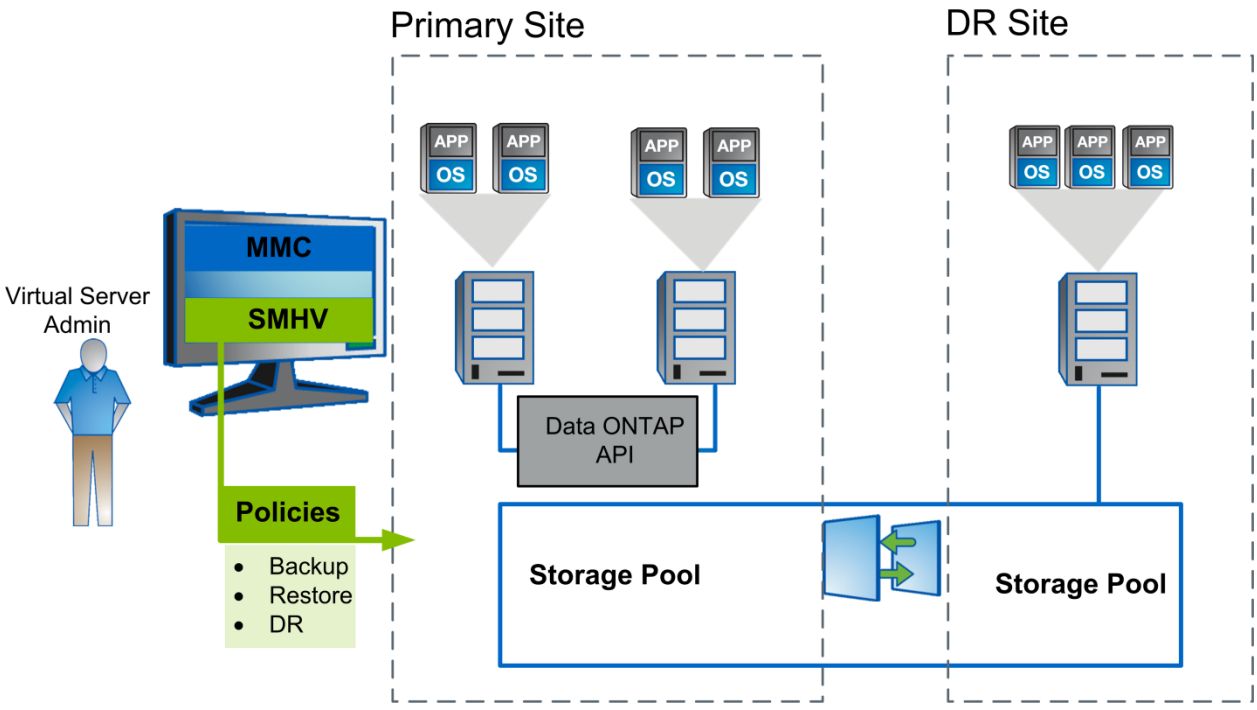

With clustered Data ONTAP systems, customers can host all VMs that pertain to a workload within volumes mapped to a NetApp Storage Virtual Machine (SVM; formerly known as Vserver). This simplifies provisioning and protection management. Figure 2 shows how SMHV can be deployed to manage virtual entities in a SAN environment with clustered Data ONTAP 8.2.

Figure 2) SMHV deployed to manage virtual entities in a clustered Data ONTAP environment.

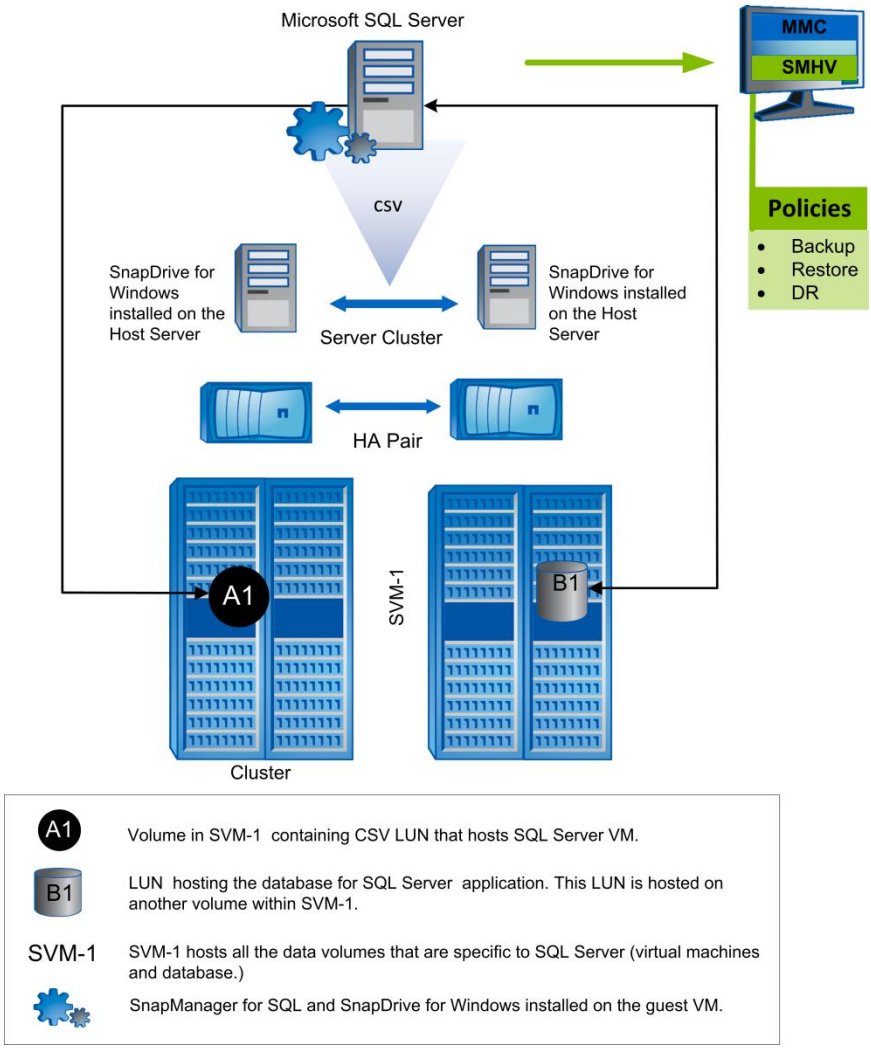

SMHV Components

SMHV SnapInfo Directory

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 and CSV LUNs only with Windows Server 2012 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.

Figure 3) SnapInfo settings.

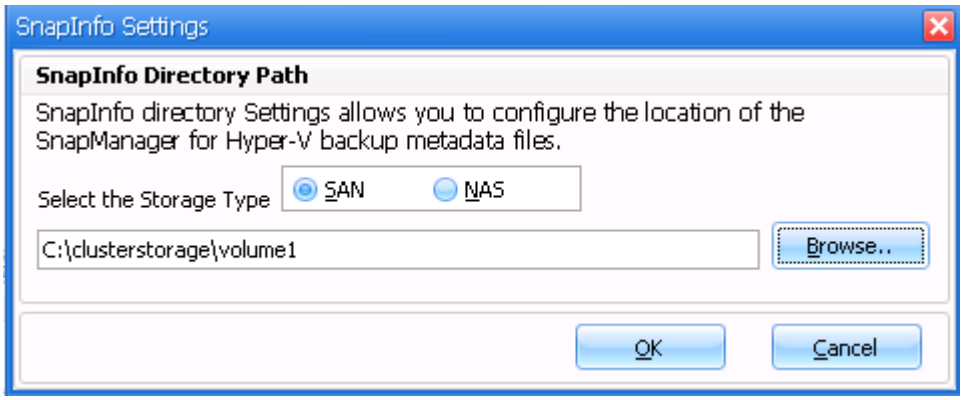

Best Practice

For clustered Data ONTAP 8.2, NetApp recommends having the SnapInfo directory or a share in a separate volume in the Storage Virtual Machine and not as part of the other data volumes. For example, if SMHV is protecting SQL Server VMs that are hosted in a CSV LUN in a volume, then the user must make sure that the SnapInfo directory is not part of this volume of the SVM. This simplifies VM restoration and disaster recovery.

SMHV Report Settings

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

Note: The report path must not reside on a CSV.

SMHV Event Notifications

The event notification setting can be configured to send e-mail and AutoSupport™ messages when an event occurs.

6 SnapManager for Hyper-V Backup Types

6.1 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.

6.2 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.

Crash-consistent backups are also supported both for SAN and SMB 3.0 environments.

Figure 4 shows the Backup Dataset Wizard showing the application-consistent and crash-consistent backup types.

Figure 4) Backup Dataset Wizard screen.

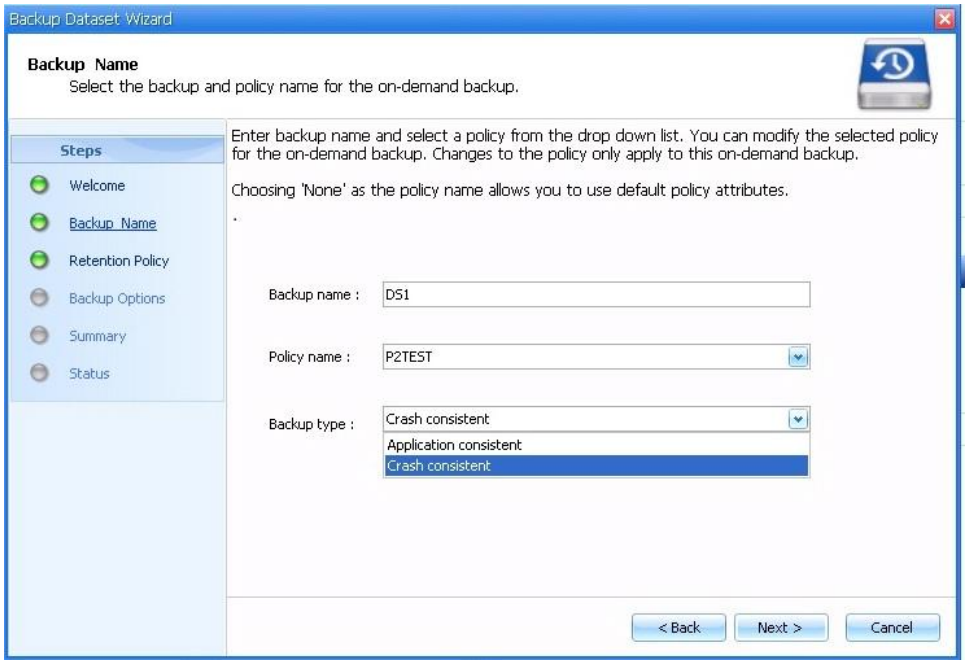

Note: Saved state backup policy is not applicable for crash-consistent backup and restore. This is because crash-consistent backups do not involve the Hyper-V VSS writer.

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.

Note: Restore of crash-consistent backups in SMB environments fails if the directories hosting them are renamed after the backup is performed.

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.

7 SnapManager for Hyper-V Process Flow

7.1 SMHV 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.

7.2 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.

7.3 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.

Figure 5 is 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 5) Hyper-V infrastructure and associated storage during an application-consistent SMHV backup.

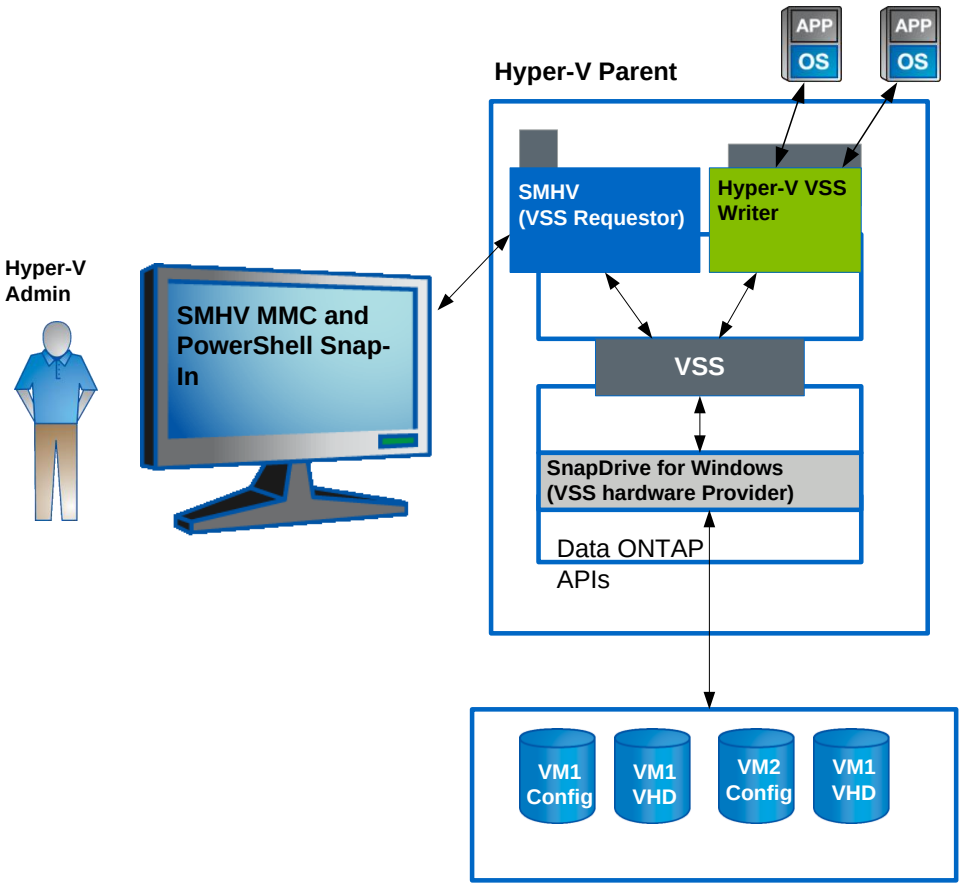

The following steps describe the flow of the backup process in a SAN environment:

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 or SnapVault relationship on the volume if the Update SnapMirror or Update SnapVault option was selected. SnapMirror and SnapVault are discussed in detail in section 9, “SnapManager for Hyper-V Disaster Recovery for SAN Environments.”

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.

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.

Best Practice

For clustered Data ONTAP systems, all VMs related to an application can be isolated to a single NetApp Storage Virtual Machine. A single dataset can be created to back up and restore these VMs. This simplifies the backup, restore, and DR processes.

7.4 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 6 illustrates the SMHV backup process for Windows Server 2012.

Figure 6) SMHV backup process for Windows Server 2012 SAN environments.

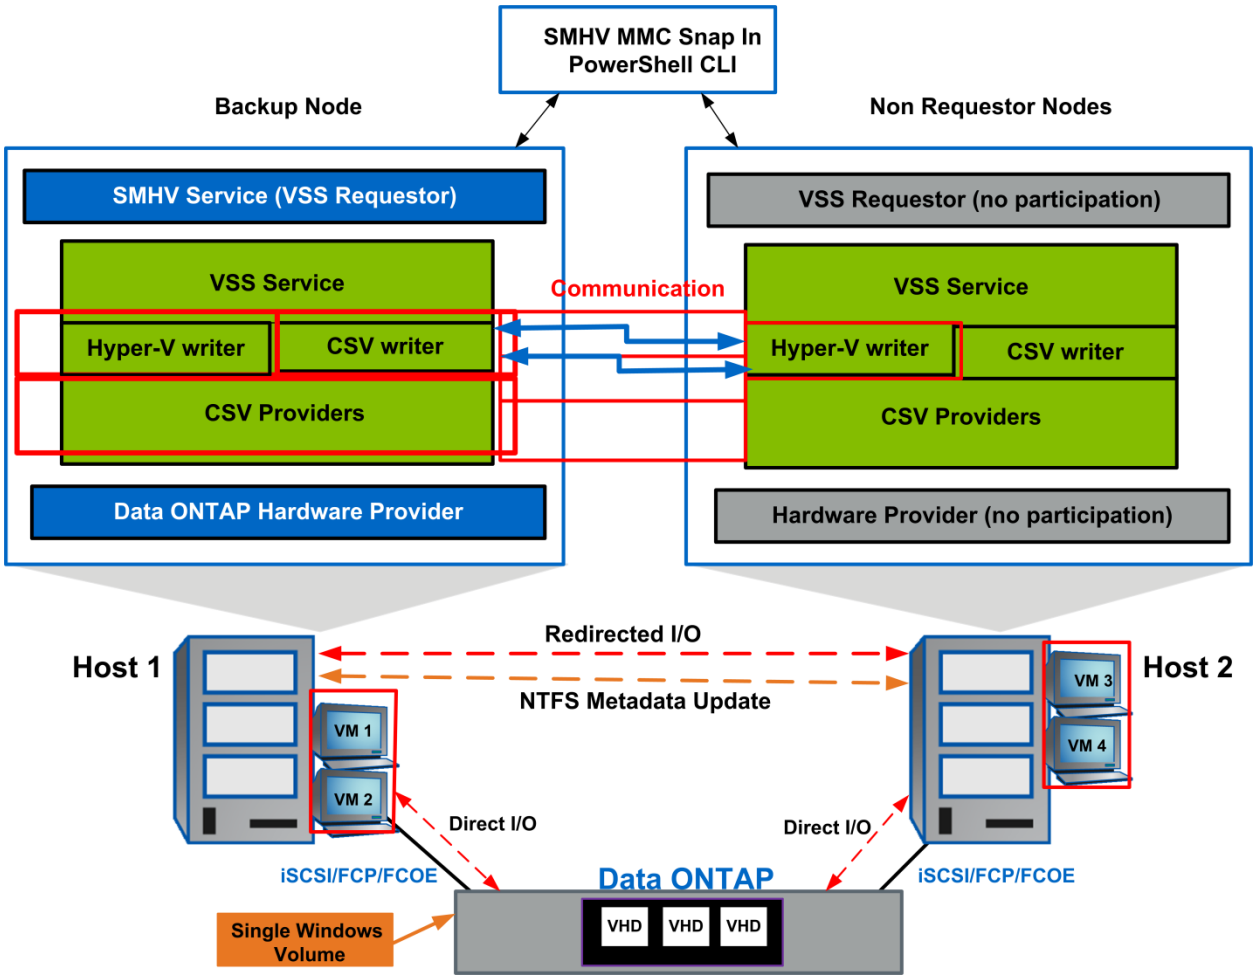

Initialization Phase

- 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 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 via virtual Fibre Channel in the VM. However, during the VM backup, the LUN presented to the VM is not backed up.

Note: If a Hyper-V cluster environment has more than 2100 VMs, crash-consistent and distributed application-consistent backups might fail.

7.5 SMHV 2.0 Backup Process in Windows Server 2012 SMB 3.0 Environments

Data ONTAP 8.2 supports two important features developed specifically for Windows Server 2012 environments: continuously available shares for Hyper-V over SMB; and remote VSS. Users can create continuously available SMB shares (by using the provisioning templates in SnapDrive 7.0 for Windows) and host virtual machines on them. These virtual machines can be backed up by using SnapManager for Hyper-V using remote VSS. Users can initiate a backup from either the SMHV GUI or PowerShell cmdlets.

When the VSS requestor (SnapManager for Hyper-V) adds an SMB 3.0 share containing Hyper-V VMs to the VSS Snapshot set, VSS invokes the new SMB File Share Copy Provider component in Windows Server 2012 to send the MSRPC commands to the SMB target (storage system) to coordinate the VSS backups. The new File Share Shadow Copy Agent (remote VSS provider) running on the SMB target is responsible for creating the actual hardware Snapshot copy. Data ONTAP 8.2 implements the File Share Shadow Copy Agent (remote VSS hardware provider) to perform the application-consistent backup copy of the SMB shares.

Figure 7) SnapManager 2.0 for Hyper-V architecture.

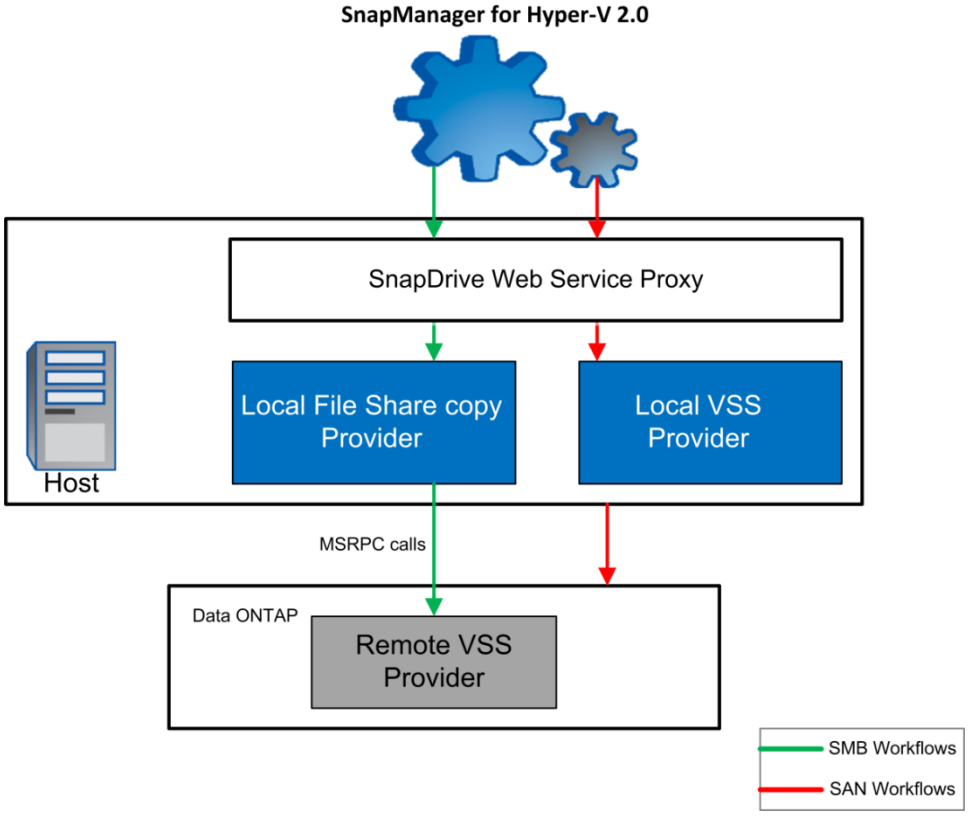

Initialization Phase

The backup mechanism for virtual machines in SMB 3.0 shares is different from that of VMs hosted on CSV LUNs. In the case of VMs on SMB shares that are part of a dataset and are spread across a cluster, the VSS backup is initiated on every node that owns the VMs. This is different from the way backups are performed on CSV LUNs. In the case of Windows Server 2012 CSV LUNs, the backup is initiated by the cluster owner node, which coordinates with the rest of the nodes to initiate the VM backups. The entire backup is executed from the cluster node only (distributed backup).

Prebackup Phase

The Hyper-V writer on the VM owner nodes quiesces the application writers inside the VM by using the integration services. This makes sure that the VM state is consistent before starting the Snapshot copy. The Hyper-V VSS writer in the parent OS coordinates backup with the integration services in the virtual machines.

Backup Phase

1. The local VSS hardware provider on each Hyper-V node initiates the backup Snapshot copy process of the VM.
2. SMHV calls a SnapDrive web service proxy to add the SMB shares that host VMs in the dataset.
3. The Hyper-V VSS writer on each node gathers the VM metadata and communicates with remote VSS on the NetApp storage system to create a backup Snapshot copy of the VM.
4. SMHV retrieves the VSS backup metadata and backup component and saves them in the SMHV SnapInfo directory.
5. Finally, SMHV registers the backup for the VM.

Note: The VSS framework does not allow granular backups for shares. This means that all the folders in the share get backed up.

Postbackup Phase

1. SMHV saves the backup metadata to the SnapInfo directory.
2. VSS renames the snapshot GUID to the SMHV naming conventions [`dataset_hostname_timestamp_backup`] for each volume.
3. Applicable policy processing is performed; for example, retention of older backups, SnapMirror or SnapVault update, execution of any specified postscript, or generation of ASUP notifications.
4. Creates a Snapshot copy of the SnapInfo directory at the volume level.
5. Applies retention and deletes older SnapInfo Snapshot copies.

Note: The SnapMirror update is done once per cluster after all the node level backups have been created.

Note: For Hyper-V over SMB, the CIFS server and SVM should have different names.

Note: SMHV 2.0 uses clones directly from the active file system during a backup, which increases the backup performance. However, cloning from AFS consumes twice as much space. This can be disabled by creating the "`DisableCloneFromAfs`" registry key on the host system located under HKEY_LOCAL_MACHINE > SYSTEM>Services > SWSvc > Parameters, and setting its value to `1`.

Note: Application-consistent backup for VMs in SMB shares is not supported with Hyper-V Server 2012.

7.6 Scheduled Backups and Retention Policies

SMHV allows administrators to schedule a dataset backup at a particular time. SMHV uses the Windows Tasks Scheduler to create or modify scheduling policies. The limit of 255 NetApp Snapshot copies per volume 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 in order not to exceed the maximum number of Snapshot copies per volume. If the number of Snapshot copies exceeds 255 on any given volume, future backups against that volume will fail.

Best Practice

Because the SnapManager suite of products (SnapManager for SQL Server, SnapManager for SharePoint, SnapManager for Exchange, and SnapManager for Hyper-V) use SnapDrive for application-consistent Snapshot copies, NetApp recommends having minimal overlaps when these application-specific Snapshot copies are initiated through their respective products. This approach reduces the performance overhead on the cluster SVM.

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 7-day period
- **Monthly.** A specified day and time within a calendar month
- **Unlimited.** Never-deleted backups

Within the selected retention type, there is a choice between deleting backups that are older than a specified period of time or deleting backups that exceed a maximum total.

NetApp recommends using the retention policies to meet specific SLAs, and 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 backup limits the maximum number of Snapshot copies associated with the backup to 60. However, if the retention policy is configured as 30 days, the Snapshot copy limit per volume is reached in 5 days, and backups fail from that point on.

Best Practice

Select a backup retention level based on the 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, there is a higher risk of not having a usable backup from which to restore in case of a disaster.

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

In addition, the NetApp Snapshot copies associated with on-demand backups must be considered when determining the number of Snapshot copies maintained against a volume.

7.7 Handling Saved-State Backups 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 must still proceed, even if an online backup is not possible. To make such backups possible, 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: Selecting the Allow Saved-State VM Backup checkbox enables SMHV to back up the VM using the saved state. If this option is selected, SMHV does 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. Performing a saved-state or offline backup can cause downtime. For more information about online and offline VM backups, refer to Planning for Backup on the Microsoft TechNet site.

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, because the VM is being backed up regardless of the state.

7.8 Backup Scripts

Postscript Variables

SMHV can be used to run optional backup scripts either before or after the backup takes place. These scripts run on all dataset member hosts unless a specific server is indicated. 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 directory at the end of the dataset backup.

During the postscript 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 VMs only.

PowerShell Cmdlets

In addition to performing common tasks through the SMHV GUI, SMHV offers PowerShell cmdlets that can be run to perform common tasks. The following PowerShell cmdlets are available when SMHV is installed:

- Add-SMHVDataSet
- Add-SMHVHost
- Add-SMHVPolicy
- Delete-Backup
- Get-Backup
- Get-SMHVDataSet
- Get-SMHVHost
- Get-SMHVPolicy
- Get-VMsFromBackup
- Invoke-SMHVConfigureHost
- Invoke-SMHVRemoteHostInstall

- Invoke-SMHVRemoteHostUninstall
- New-Backup
- Remove-SMHVDataSet
- Remove-SMHVHost
- Remove-SMHVPolicy
- Restore-Backup
- Set-SMHVDataSet
- Set-SMHVPolicy

Use the Help option to learn how to use these cmdlets.

7.9 Quick and Live Migration Best Practices

- SMHV cannot back up a VM that is actively undergoing migration. When a backup runs against a dataset in which VMs are actively being migrated, an error is generated, and those particular VMs are not backed up. It is best to avoid SMHV-related operations in the virtual machine during live migration.
- It is best to avoid SMHV-related operations during storage live migration. Otherwise, such operations could corrupt the virtual machine. After performing storage migration of a VM from one LUN or share to another LUN or share, the virtual machine can no longer be restored with the previous Snapshot copy. As a safety net, create an SMHV backup immediately after VM storage migration is complete.
- After performing storage migration of a VM from one share to another share hosted on a different volume, the SnapMirror and SnapVault relationship must be verified. Also, previously existing Snapshot copies cannot be restored from the SnapVault storage system.

7.10 Restore Process

SMHV can restore a virtual machine from a backup. It can also restore a VM that is part of a cluster. To restore the VM, SMHV uses the file-level restore feature in SnapDrive for Windows. The associated files of a VM, including the configuration file, Snapshot copies, and any VHDs, can be spread across various Data ONTAP LUNs. A LUN can contain files belonging to various VMs.

Note: Make sure that the SnapRestore license is present on the storage system before attempting to restore.

If a LUN contains only files associated with the VM to be restored, SMHV restores the LUN by using LUN clone split restore. If a LUN contains files not associated with the VM that is to be restored, SMHV restores the VM by using the file-level restore operation (SIS clones).

With these differences in restore types aside, the SMHV uses the following process flow during a restore:

1. SMHV restores a VM in coordination with Hyper-V VSS writer. Hyper-V VSS writer powers off the VM and deletes it before the 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 is added back into 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:

5. The VM to be restored is not [currently running] on the host.
6. The 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.
7. The Snapshot location of the VM has changed.

8. 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 first the user must confirm that the restore should proceed.

   Note: If the VM no longer exists, it can still be restored 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, it can still be restored by selecting a backup to which it belonged.

   If the VM was removed from all datasets before it was deleted, it can still be restored 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 the same, the restored VM might have additional or fewer VHDs. If that is the case, NetApp recommends manually updating the cluster configuration of the VM and its dependencies.

Best Practice

In the event of a VM restore failure and VM deletion, NetApp recommends manually copying the virtual machine data and adding the following registry settings, using the command line on the host where VM restore is being attempted:

```
REG ADD HKLM\SYSTEM\CurrentControlSet\services\SnapMgrServiceHost\Parameters
/v RestoreVMFromExistingData /t REG_DWORD /d 1
```

After this, VM restore can be attempted using the SMHV GUI.

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, the VM can be restored from SMHV, but it must be manually made highly available. For more information, refer to the Failover Clustering on Windows Server 2008 R2 article.

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 in Windows 2008 R2. This is supported in Windows server 2012.

Note: If a crash-consistent backup is created by the earlier versions of SMHV, it cannot be restored by using the latest version of SMHV.

7.11 Mounting a Backup

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

To mount a backup, perform these steps:

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

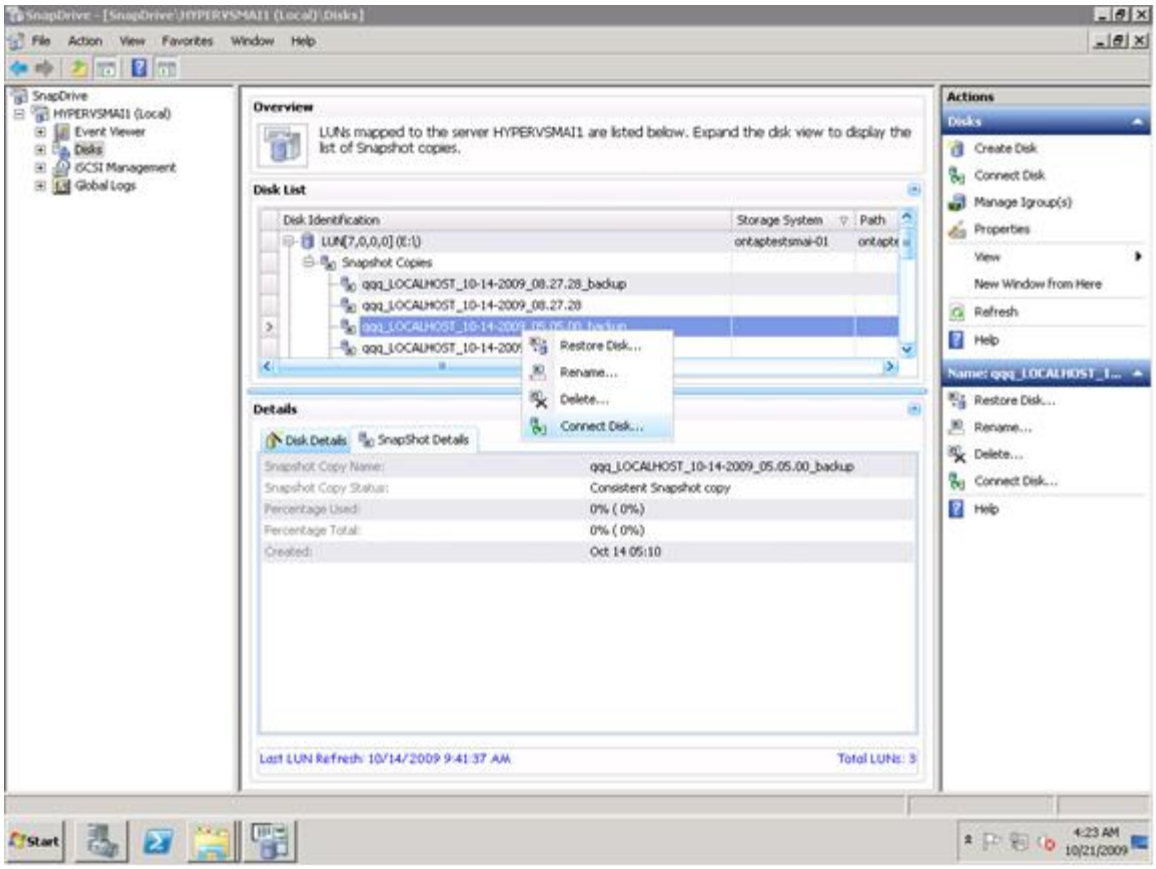

2. Right-click the Snapshot copy (the one with the `_backup` suffix) and select the Connect Disk option.
3. Click Next.
4. If the LUN is a dedicated disk, proceed to step 5. If the LUN is a Windows cluster resource, perform the following step in the Specify Microsoft Cluster Services Group panel. In the panel, select only one of the following actions and then click Next:
   - To use an existing cluster group, select it from the Group Name drop-down list.
   - To create a new cluster group, select the Create a New Cluster Group option.

Note: When selecting a cluster group for the LUNs, choose the cluster group the application will use. If a volume mount point is being created, the cluster group is already selected. This is because the cluster group owns the root volume physical disk cluster resources. NetApp recommends creating new shared LUNs outside of the cluster group.

5. To use CSVs, select the Add option.
6. In the Select LUN Properties panel, either select a drive from the list of available drive letters or enter a mount point for the LUN that is being connecting. When a volume mount point is created, enter the drive path that the mounted drive will use (for example, `G:\mount_drive1\`).
7. In the Select Initiators panel, choose an initiator for the LUN.
8. In the Select Initiator Group Management panel, specify either automatic or manual igroup management.
9. In the Completing the Connect Disk Wizard, perform the following actions:
   a. Verify all of the settings.
   b. To change any settings, click Back to return to the previous wizard panels.

c. Click Finish.

10. Browse the backup by selecting the drive letter on the explorer of the 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. Once a backup has been mounted, the Hot Disk Add functionality in Windows Server 2008 R2 can be used to attach a disk (backed by the VHD) to the VM at run time without shutting down the VM. This functionality makes it possible to attach new disks to the VM.

This is a three-step process:

1. Mount the VHD from the backup mounted location (`<drive>:\ Name.vhd`) to the parent host by using the Attach VHD option from the Disk Manager snap-in. This mounts the VHD as a new disk in the Hyper-V parent.
2. Offline the disk mounted in step 1 by using the Disk Manager snap-in. Select the disk and select the Offline menu item. This offlines the disk that was mounted from the VHD.
3. Attach the offlined disk to the VM by clicking the Physical Hard Disk button. Select the offlined disk from the Physical Hard Disk drop-down list. This presents a new drive inside the VM (backed by VHD in the parent).

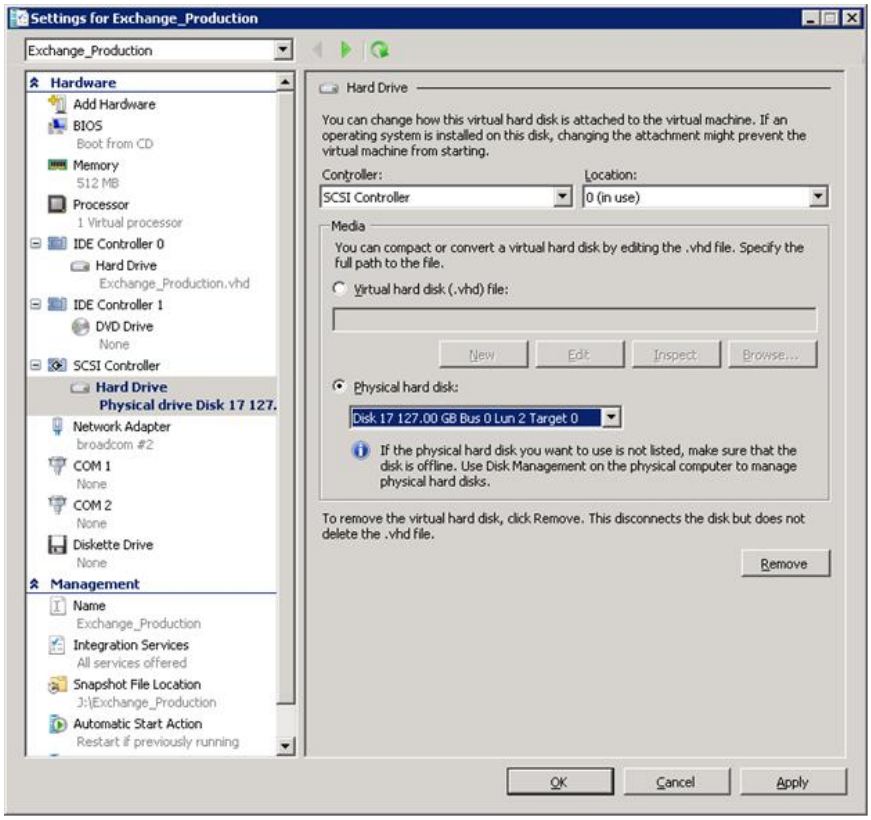

4. Log into the VM and select the newly mounted drive to see the contents of the disk backed by the VHD attached.
5. When verification is complete, detach the disk from the VM, using the VM settings screen, and click Remove. Use SDW to unmount the disk backed by the Snapshot copy, using the SnapDrive

Disconnect disk MMC action/menu item. Alternatively, use the `SDCLI Snap_Unmount` command to unmount the disk mounted from the Snapshot copy.

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.

8 SnapManager for Hyper-V High Availability

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

Note that SMHV, as a host-end application, offers services provided that the storage is continuously available. Section 8.1 describes the available tools that facilitate storage availability.

8.1 Multipath High Availability 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 Fibre Channel (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 an active-active storage controller configuration to eliminate any single points of failure (SPOFs).
- Use multipath HA with an active-active controller configuration to improve storage availability and performance.

For more details on HA system configuration, refer to NetApp TR-3450: Active-Active Controller Overview and Best Practices Guidelines.

8.2 Nondisruptive Operations

To accommodate storage failover, two storage controllers (nodes) in the same cluster are connected together as an SFO pair, called an “active-active” pair. Each node of the pair is a fully functioning node in the cluster (hence the term “active-active”). Clusters can be heterogeneous (in terms of hardware and Data ONTAP versions), but an SFO pair must be the same controller model.

Asymmetric Logical Unit Access (ALUA) support is available in all clustered Data ONTAP configurations with Data ONTAP DSM as well as with MSDSM. Hosts use ALUA to determine the state of a specified path during a failover.

This feature enables virtualized workloads to fail over nondisruptively to another storage system.

Note: During a storage failover, any activity related to SnapDrive and SnapManager fails. All operations continue to work after the failover is complete.

Figure 8) Storage failover.

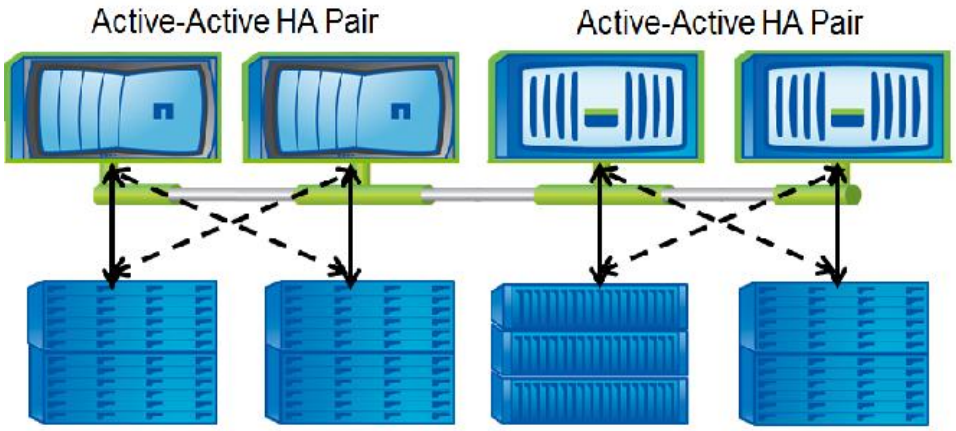

Clustered Data ONTAP supports volume moves from one aggregate to another aggregate anywhere within the cluster. Volume moves are performed when cluster expansion is needed; when the business priority changes; or as part of a hardware lifecycle management task.

A volume is moved by using a single administrator command, from the CLI or System Manager. Volume moves are nondisruptive and provide uninterrupted access to the hosts that are connected to the storage using any protocol—iSCSI, NFS, CIFS, FC, or FCoE—during the volume move.

This means that users can continue to access data while a volume move is being performed in the background. All sessions from SnapDrive to the storage systems continue to exist.

Note: SnapDrive and SnapManager are transparent to the volumes that are currently in a migration state. All volumes are displayed in the Create/Connect Wizard, regardless of their state.

Note: The cutover window defined for a SAN volume should not exceed the expected timeout value on the host side. During the cutover phase of the volume move, all I/O access is queued, and requests are blocked to the source volume. SnapDrive sets a timeout value of 120 seconds on the host during the volume move. In addition, when a SAN volume is moved, ALUA is used to optimize access to the volume. NetApp recommends performing volume moves during nonpeak hours.

Note: Each node must have a data LIF for optimized access to the volume.

Best Practices
• Use an active-active storage controller configuration to eliminate any SPOFs. • Use multipath HA with an active-active storage configuration to get better storage availability and higher performance. • For more details on HA system configuration, refer to TR-3450: Active-Active Controller Overview and Best Practices Guidelines.

9 SnapManager for Hyper-V Disaster Recovery for SAN Environments

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

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
```

9.1 Get-VMsFromBackup Cmdlet

This new 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 and must know which VMs are present in the backup in order to restore them on the secondary. This 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
```

9.2 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).
- 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 directory.

9.3 To Fail Over VMs to the Secondary Site

Following are the steps to fail over VMs to secondary:

1. Connect to all of the LUNs from secondary storage system volumes. If the secondary is a cluster, go to the node where a cluster group is on line 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.
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 is a cluster, make sure that the LUNs on which VMs reside are online on the cluster node that owns the cluster group.
7. If the secondary 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 above example, a VM named “testVM” having four VHDs is stored in primary host in J: drive .It is restored to the K: drive in the secondary. K:\ is mapped to secondary SnapMirror volume.

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

1. Connect the primary SnapInfo LUN to the secondary host by breaking the mirrored volume.
2. Snap restore from the last SMHV Snapinfo Snapshot copy.
3. Copy the contents to the already existing SnapInfo copy to the secondary.

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

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

1. Connect the primary SnapInfo LUN to the secondary host by breaking the mirrored volume.
2. Use SnapRestore to restore from the last SMHV SnapInfo Snapshot copy.
3. Copy the contents to the already existing SnapInfo copy to the secondary.

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

9.4 To Fail Back VMs to the Primary Site

Follow these steps to fail back VMs to the primary site:

1. Get the data from the secondary 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 secondary to primary (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 down temporarily, the user must get to the primary only those changes that happened on the secondary while the primary was gone. To do this, resync the existing SnapMirror relationship in the reverse direction (resync from the secondary to the primary).
4. When the data on the secondary is synchronized with the primary, go to the SnapDrive user interface on the secondary and initiate a SnapMirror update for each of the LUNs on the secondary. If this is not done, SDW uses the SMHV backup Snapshot copy to restore the LUNs on the primary during the resyncing in step 3. The LUN in the backup Snapshot copy is actually a LUN clone, so this must be avoided by forcing one more SnapMirror update.

Note: Creating an SMHV backup (with the Update SnapMirror option checked) from the secondary has the same effect as manually doing the SnapMirror update from the SDW GUI. Most users will take the SMHV backup instead 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 (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 to the primary. This should be released.

6. Restore the SnapInfo directory 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 8) and backup (from step 9). Use the `-RestoreToAlternateHost` switch and specify the primary host or cluster name as `-server` parameter. If the primary 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 is the cluster, make the VMs highly available by using the failover cluster UI or Windows PowerShell cmdlet.

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

1. If the secondary is a standalone host, shut down and delete the VMs running on it. Disconnect the SnapInfo directory and the disks containing VMs that use SnapDrive. If the secondary 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 that use SnapDrive.
2. Resync the SnapMirror relationship from the primary to the secondary.

10 SMHV Disaster Recovery for VMs in Hyper-V Over SMB Environments

The restore-backup cmdlet can also be used to restore VMs in SMB environments in the secondary site. This can be achieved by using the restore-backup cmdlet with the following parameters:

- RestoreToAlternateHost
- VirtualMachinePath
- SnapshotFilePath
- VHDs

10.1 Prerequistes

- 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).
- 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 directories.
- Make sure that the access control lists on the share are set appropriately. If this not done, errors such as “access denied” will appear. Refer to the section "Configuring an applying file security on NTFS files and folders" in the Clustered Data ONTAP 8.2. File Access Management Guide for CIFS.
- The virtual switch name for a VM must be exactly the same for the primary and secondary Windows host

10.2 Steps

1. Make sure that when the `Get-VMsFromBackup` cmdlet is run on the hosts in the secondary site, it lists all the VMs in the share as shown below.

```
PS C:\> Get-VMsFromBackup

Name                                                    Id
----                                                    --
vm2013                                                  224DDECE-5C50-401F-8B48-9797EBD58CC6
DR-VHD2                                                 582E6C41-DC55-4798-A48A-91B4930C1224
DR-SS                                                   0DE34644-CA67-4772-8E91-F14A7584966C
DR-VHD2-2SHARES                                         F4D63D8B-4F59-4C99-AC99-E1D3F94C0914
```

2. Obtain the backup information by using the `get-backup` cmdlet.

```
PS C:\Users\administrator.SDDEV> Get-Backup -Resourcename vm2013


BackupName     : ds-new_01-14-2013_14.20.48
FilesList      : {\\172.17.175.81\vol2_share\vm2013\Virtual Machines\224DDECE-5C50-401F-8B48-9797EBD58CC6.xml,
                 \\172.17.175.81\vol2_share\vm2013\Virtual Machines\224DDECE-5C50-401F-8B48-9797EBD58CC6\*,
                 \\172.17.175.81\vol2_share\vm2013\Virtual Hard Disks\vm2013.vhdx}
RetentionType  : hourly
DatasetName    : ds-new
BackupId       : ds-new_01-14-2013_14.20.48
BackupTime     : 1/14/2013 2:20:48 PM
BackupType     : Application consistent
BackedupVMs    : {vm2013}
```

3. Perform restore by using the `restore-backup` cmdlet as shown.

```
PS C:\Users\administrator.SDDEV> Restore-Backup -server clab-a13-15 -ResourceId 224DDECE-5C50-401F-8B48-9797EBD58CC6 -Ba
ckupName ds-new_01-14-2013_14.20.48 -RestoreToAlternateHost -Verbose -DisableVerifySnapshot -VirtualMachinePath "\\172.1
7.175.81\vol5_share\vm2013" -SnapshotFilePath "\\172.17.175.81\vol5_share\vm2013" -VHDs @(@{"SourceFilePath" = "\\172.17
.175.81\vol2_share\vm2013\Virtual Hard Disks\vm2013.vhdx"; "DestinationFilePath" = "\\172.17.175.81\vol5_share\vm2013\Vi
rtual Hard Disks\vm2013.vhdx"})_
```

11 SnapVault Integration

SMHV 2.0 supports SnapVault in clustered Data ONTAP 8.2 environments. After SMHV completes a backup operation for a dataset, the backup can be updated to the SnapVault destination.

Figure 9) SnapVault integration in SMHV.

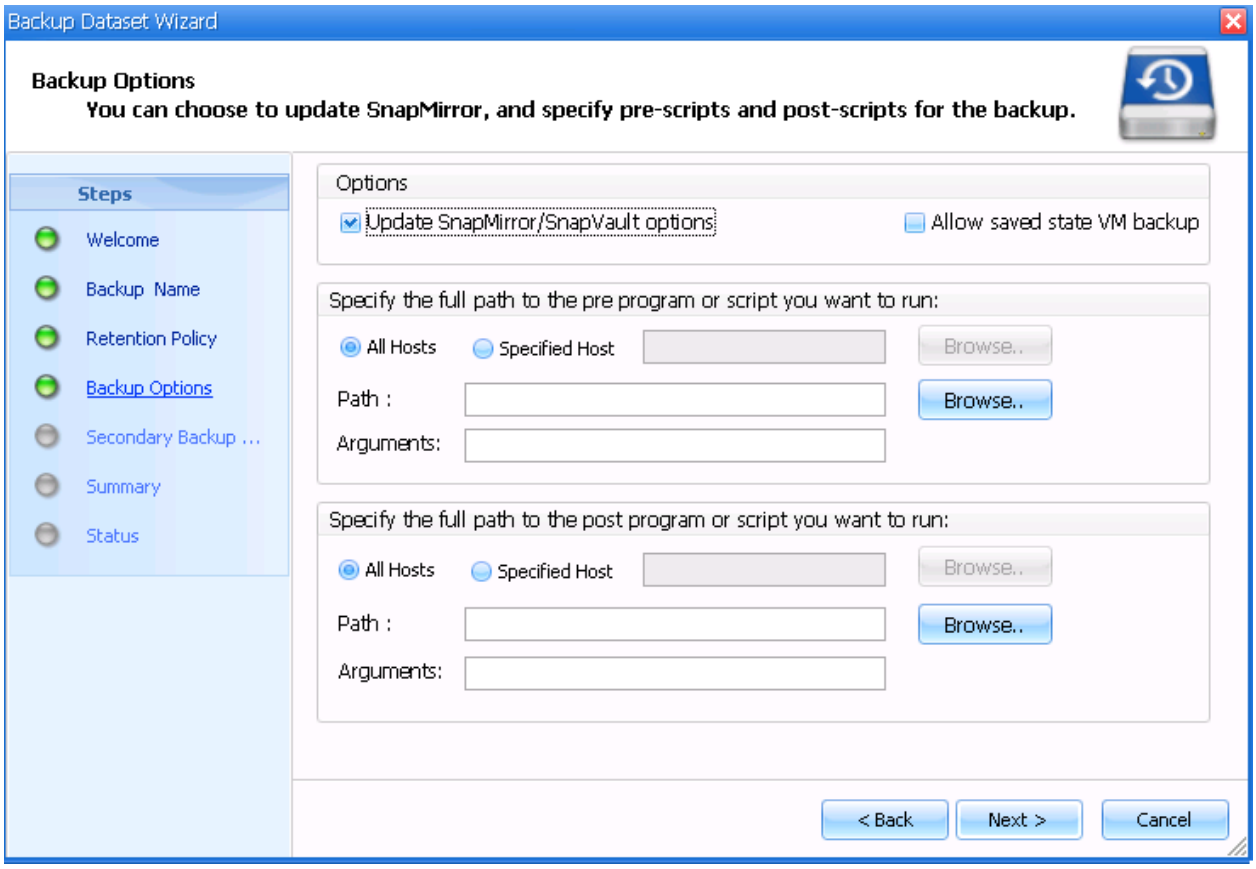

You can also configure the retention period of a Snapshot in the SnapVault storage system. This can be done by selecting the labels option while configuring the SMHV backup workflow as shown in Figure 10.

Figure 10) SMHV SnapVault options and Snapshot labels.

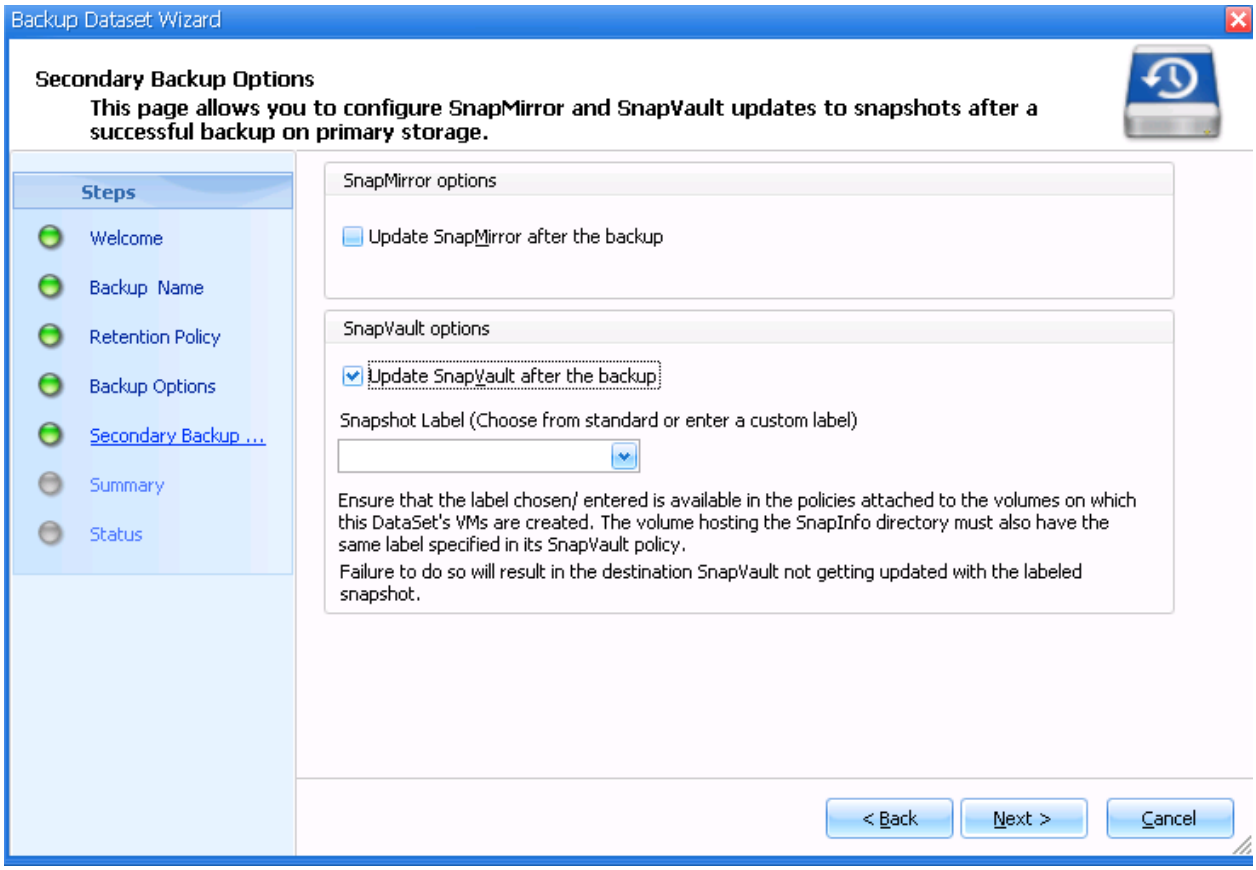

You can choose from a set of preconfigured labels such as `smhv_hourly`, `smhv_daily`, `smhv_weekly`, `smhv_monthly` or can choose to have a custom label. This is required to be performed on the storage system and attached to the policy of the SnapVault relationship between the primary and the secondary volumes.

After the dataset backup is created, the primary backup is labeled with the Snapshot label and updated to the SnapVault destination using SnapMirror updates.

Each Snapshot update to SnapVault destination will be tied to a Version UUID. This Version UUID along with other information such as SnapVault status and label details are stored in the Snap Info metadata.

Note: SnapVault restore using SMHV is not supported.

Best Practice

Before initiating backup and SnapVault operations using SMHV, make sure that the underlying volumes for all VMs belonging to a dataset have the same SnapMirror label and SnapVault policies.

12 Scalability and Performance

12.1 Scaling beyond 1000 VMs

If your Hyper-V host or host cluster has more than 1000 virtual machines, you must increase the value of the `maximumElementsInCacheBeforeScavenging` property in `SnapMgrServiceHost.exe.config` file for Hyper-V Cache Manager. This value should be greater

than or equal to the number of Hyper-V hosts running on a standalone host or cluster. The value should be changed on each node of the cluster. Restart SnapManager for Hyper-V service after changing this value. You must manually edit the `SnapMgrServiceHost.exe.config` file using a text editor.

12.2 VSS Limitation

VSS requires that the provider commit 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.

12.3 Quality of Service with Data ONTAP 8.2

Storage QoS is a new feature in Data ONTAP 8.2 that provides the ability to group storage objects and set throughput limits on the group. With this ability, a storage administrator can separate workloads by organization, application, business unit, or production or development environments.

In enterprise environments, storage QoS helps to achieve the following:

- Prevents user workloads from affecting each other
- Protects critical applications that have specific response times that must be met In IT as a service (ITaaS) environments, storage QoS
- Prevents tenants from affecting each other
- Avoids performance degradation with the addition of each new tenant

QoS allows you to limit the amount of I/O sent to a Storage Virtual Machine, a flexible volume, a LUN, or a file. I/O can be limited by the number of operations or the raw throughput.

Figure 11) SVM with its own QoS policy.

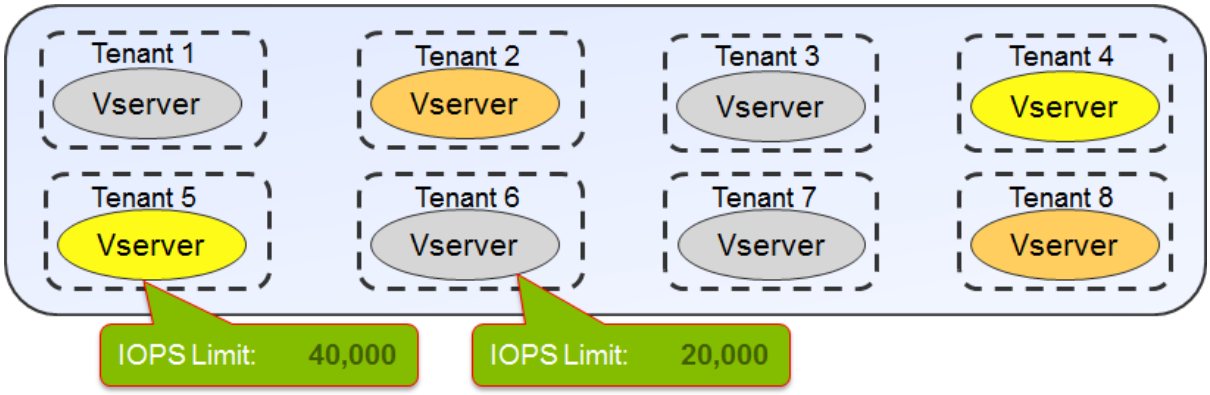

Storage side QOS feature can be used as a solution for complying with service levels at the storage layer in cloud environments. Virtual machines workloads can be isolated to an SVM and tied to a customer. QoS policies limiting the IOPS can then be assigned to these Vservers to comply with service-level agreements. All VMs belonging to one customer can be grouped into dataset in SMHV and assigned a backup policy. By assigning a QOS policy in conjunction with backup, retention, and replication policies in SMHV, cloud service providers can make sure that service-level agreements are met at the storage layer.

13 SnapManager for Hyper-V Conclusion

SnapManager for Hyper-V offers 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 can help administrators to get the most out of SMHV deployments.

Appendixes

Clustered Data ONTAP 8 Terminology

Table 4 describes the terms used in clustered Data ONTAP 8.2.

Table 4) Terminology used in clustered Data ONTAP 8.2.

Term	Description
cluster	In clustered Data ONTAP, a group of connected nodes (storage systems) that share a global namespace and that can be managed as a single virtual server or multiple virtual servers, providing performance, reliability, and scalability benefits. In the Data ONTAP 7.1 release family and earlier releases, a pair of storage systems (sometimes called nodes) configured to serve data for each other if one of the two systems stops functioning.
cluster interconnect	A 10GbE network connection for data communication across nodes.
high availability (HA)	In Data ONTAP, the recovery capability provided by a pair of nodes (storage systems), called an HA pair, that are configured to serve data for each other if one of the two nodes stops functioning.
HA pair	In Data ONTAP, a pair of nodes (storage systems) configured to serve data for each other if one of the two nodes stops functioning.
intracluster replication	SnapMirror replication across NetApp Storage Virtual Machines (formerly Vservers) residing in two different clustered Data ONTAP systems.
LIF	A logical network interface, representing a network access point to a node. LIFs currently correspond to IP addresses, but they can be implemented by any interconnect. A LIF is generally bound to a physical network port; that is, an Ethernet port. LIFs can fail over to other physical ports (potentially on other nodes) based on policies interpreted by the LIF manager.
node	A storage controller (one storage controller equals one node; an HA pair equals two nodes).
Storage Virtual Machine (SVM. Formerly known as Vserver (virtual server).	A secure virtual storage server that supports multiple protocols and unified storage: • Contains data volumes and one or more LIFs through which it serves data to the clients. • Securely isolates the shared virtualized data storage and network and appears as a single dedicated server to the clients. Each SVM has a separate administrator authentication domain and can be managed independently by an SVM administrator. • Represents a single file system and a unique global namespace. A global namespace enables the NAS clients to access data without specifying the physical location of the data. The global namespace also enables cluster and SVM administrators to manage distributed data storage as a single file system.

To Deploy Clustered Data ONTAP Storage System

Follow these steps to deploy a storage system that uses clustered Data ONTAP 8.2

1. Set up the cluster environment (refer to the Data ONTAP 8.2 Installation and Administration Guide).
2. Create an aggregate.
3. Create a Storage Virtual Machine. The following screen shows sample SVM (Vserver) properties.

```
APPCL::> vserver show -vserver infraserver

                                    Vserver: infraserver
                                Vserver Type: cluster
                                Vserver UUID: d3aa46e2-97f6-11e0-bbc3-123478563412
                                Root Volume: infraroot
                                  Aggregate: infraggr
                        Name Service Switch: file
                        Name Mapping Switch: ldap
                                  NIS Domain: -
                    Root Volume Security Style: ntfs
                                 LDAP Client: -
                                    Language: C
                               Snapshot Policy: default
                                     Comment:
                  Anti-Virus On-Access Policy: default
                                 Quota Policy: default
                    List of Aggregates Assigned: -
Limit on Maximum Number of Volumes allowed: unlimited
                           Vserver Admin State: running
                             Allowed Protocols: nfs, cifs, fcp, iscsi
                          Disallowed Protocols: -
```

4. Create iSCSI service to set up the iSCSI target node.
5. Configure the network for the SVM:
   - The data LIFs, which enable SVMs to serve data to the clients (iSCSI, FCP, CIFS)
   - The management LIF, which allows SnapDrive to communicate with the other LIFs to serve data
6. Create data volumes of the required size. SnapDrive uses these volumes for LUN creation and management.
7. For data protection within the cluster, follow these additional steps:
   a. Create a volume in the SVM of the secondary SVM. Make sure that the property of that volume is of the data protection (DP) type.
   b. Establish a SnapMirror or SnapVault relationship between the primary and the secondary by accessing the secondary system.

   For intercluster SnapMirror replication, make sure that at least one intercluster management LIF is present in each node on both the primary and the secondary storage systems.

   For information about data protection, refer to the Data Protection Best Practices Guide.
8. After creating the data volumes in the storage system, create clustered LUNs of the desired size by using SDW 7.0.
9. Create and host the required VMs in the LUNs.

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

Making choices is an unavoidable part of the process of determining the appropriate storage container format for deploying VMs by using Hyper-V.

Table 5 summarizes the pros and cons of each choice to help make the decision-making process easier.

Table 5) Choosing the Hyper-V and VHD storage container format.

Storage Container	Pros	Cons
Pass-through disk	• Delivers fastest performance • Has simplest storage path because the file system on the host is not	• Cannot create VM Snapshot copy • Is used exclusively and directly by a single VM

Storage Container	Pros	Cons
	involved • Has better alignment under SAN • For shared storage–based pass-through, no need to mount the file system on host, which might speed up VM live migration • Has lower CPU utilization • Supports very large disks	• Cannot be backed up by the Hyper-V VSS writer or any backup program that uses the Hyper-V VSS writer
Fixed-sized VHD	• Delivers highest performance of all VHD types • Has the simplest VHD file format to provide the best I/O alignment • Has more robust than dynamic or differencing VHD because of the lack of block allocation tables (redirection layer) • Offers more management advantages than pass-through disk because of its file-based storage container • Can be expanded to increase the capacity of VHD • No risk of the underlying volume running out of space during VM operations	• Might increase storage cost because of up-front space allocation when a large of number of fixed VHDs are deployed • Requires time-consuming creation for large fixed VHD • Cannot shrink the virtual capacity (reduce the virtual size)
Dynamically expanding or differencing VHD	• Delivers good performance • Is quicker to create than fixed-sized VHD • Grows dynamically to save disk space and provide efficient storage usage • Is more nimble in transporting across the network because of smaller VHD size • Does not allocate blocks of full zeros and therefore saves space under certain circumstances • Can have compact operation to reduce physical file size	• Might have I/O alignment issues caused by interleaving of metadata and data blocks • Might cause write performance to suffer during VHD expansion • Has a limit of 2040GB • Might get VM paused or VHD yanked out if disk space is running out because of dynamic growth • Does not support shrinking the virtual capacity • Cannot expand for differencing VHDs because of the inherent size limitation of the parent disk • Is not recommended for defrag because of inherent redirectional layer

SMHV: Virtual Machine Self-Management

If a VM belongs to a host that has SMHV installed, and SMHV is installed on that VM so that it can be used as a management console, do not use SMHV to manage the host to which the VM belongs.

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

Doing this and trying to restore the VM from itself causes the VM to be deleted or restarted from Hyper-V Manager.

SMHV: Data ONTAP VSS Hardware Provider Requirement

Data ONTAP VSS hardware provider must be installed in order 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 and later and does not have to be installed separately.

To View Installed VSS Providers

To view the VSS providers installed on the host, follow these steps.

1. Select Start > Run and enter the following command to open a Windows command prompt: `cmd`.
2. At the prompt, enter the following command:

```
Vssadminlist providers
```

3. 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
Provider Id: {ddd3d232-a96f-4ac5-8f7b-250fd91fd102}
Version: 6.2.0.xxxx
```

To Verify That the 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 task:

In MMC, select System Tools > Event Viewer > Application and look for an event with the following values:

```
Source Event ID Description
The VSS provider has successfully completed CommitSnapshots for SnapshotSetId 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 might be exceeded because of a transient problem. If this event is logged for a failed backup, retry the backup.

SMHV: If Virtual Machine Backups Take Too Long to Complete

If a VM contains several direct-attached iSCSI LUNs or pass-through LUNs and SDW is installed on the VM, the VM backup can take a long time. The Hyper-V writer creates a hardware Snapshot copy of all the LUNs in the VM by using the SDW VSS hardware provider. There is a Microsoft hotfix that uses the default system provider (software provider) in the VM to make the Snapshot copy. As a result, the Data ONTAP VSS hardware provider is not used for Snapshot creation inside the child OS, and the backup speed increases. For more information on the Microsoft hotfix, see Microsoft KB ID: 975354.

SMHV: Redirected I/O and Virtual Machine Design Considerations

Although redirected I/O is handled in a Windows Server 2008 R2 Hyper-V cluster, server messenger block (SMB) application programming interface (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 manually assigning 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.

To create SMHV backup datasets, follow these steps:

1. Using SDW, create one CSV per host cluster node, based on 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 by 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 reside within the same CSV.
4. Create a backup policy for each dataset that matches the appropriate 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 by using a Windows PowerShell script that can be executed by SMHV at the beginning of the backup job.

Guidelines for SMHV on Clustered Data ONTAP Systems

To avoid errors and other issues, you must make sure that the datasets and policies are scheduled so that only one backup operation is in progress on a storage system volume at any given time when using SnapManager for Hyper-V to manage systems operating in clustered Data ONTAP.

You can meet this requirement in the following ways:

- Collect all VMs that share a storage system volume in one dataset.
- Do not start application-consistent and crash-consistent backups at the same time.
- If any other applications are making Snapshot copies for the same dataset on the same storage system volumes, make sure that the time of the Snapshot operation does not overlap with the application-consistent backup schedule.

    Note: This also includes overlapping SnapMirror updates and volume move operations initiated for the volumes being backed up.

1. Make sure that VMs running across multiple dedicated Hyper-V hosts or clusters do not share a storage system volume.
2. NetApp recommends having one CSV LUN per volume and adding all the VMs belonging to this volume to the same dataset.
3. NetApp recommends segregating VMs into different CSVs based on the OS type. This will improve deduplication efficiency.
4. If the antivirus software is installed on the host, make sure that exclusions for VM files are added in the antivirus software.

References

NetApp Knowledge Base Articles

- 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: 1010887: SMHV: How to set up SnapInfo Logical Unit Number (LUN)
- KB ID: 2010607: SMHV: Creation of two Snapshot copies for every backup
- KB ID: 2010899: SMHV: Backups fail for Hyper-V virtual machines containing pass-through or iSCSI in guest disks
- KB ID: 2014900: SnapManager for Hyper-V backup sets that contain Windows XP fail
- KB ID: 2014905: SnapManager for Hyper-V backups fail to complete even though all virtual machines are located on NetApp LUNs
- 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: 3011206: SMHV: Can SnapManager 1.1 for Hyper-V exclude virtual hard disks from backups?
- KB ID: 302577: How to use the Sysprep tool to automate successful deployment of Windows XP
- KB ID: 958184: Virtual machine backup operations fail in Windows Server 2008 when Hyper-V virtual machine files are saved on a volume that is mounted on a failover cluster by using a volume GUID
- KB ID: 974909: The network connection of a running Hyper-V virtual machine is lost under heavy outgoing network traffic on a Windows Server 2008 R2-based computer
- KB ID: 975354: A Hyper-V update rollup is available for Windows Server 2008 R2
- KB ID: 975921: You may be unable to perform certain disk-related operations after an exception when a hardware provider tries to create a snapshot in Windows Server 2008 R2 or in Windows 7
- KB ID: 978157: MPIO removes pseudo-LUN paths for external storage devices on Windows Server 2008 SP2-based servers or on Windows Server 2008 R2-based servers
- KB ID: 979743: You cannot use an MPIO storage device after a failover operation in Windows Server 2008 or in Windows Server 2008 R2
- KB ID: 2406705: Some I/O requests to a storage device fail on a fault-tolerant system that is running Windows Server 2008 or Windows Server 2008 R2 when you perform a surprise removal of one path to the storage device
- KB2770917: This is a Windows Server 2012 KB fix. This is to fix the 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 Version History
Version 1.0	September 2013	Initial release

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®

www.netapp.com