Technical Report

Clustered Data ONTAP 8.1.1: Best Practices for NetApp SnapManager for Hyper-V

Santhosh Harihara Rao, NetApp
January 2013 | TR-4004

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® 8.1.1 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 and recovery 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 Virtual Machine Provisioning

Virtual infrastructure solutions such as Microsoft Hyper-V empower IT organizations to rapidly deploy VMs in all phases: development, test, and production. The tasks involved in deploying VMs usually generate many physical copies of VM images. This creates a demand for more storage resources to maintain the many VM instances and more management resources to execute the many manual steps required to deploy these VMs individually.

The integration of Microsoft virtual environments with NetApp storage technology can solve these challenges by helping organizations reduce the efforts spent deploying individual VMs and reduce the amount of storage required to support the deployment of individual VMs. It can also reduce costs associated with space, power, and cooling. The use of NetApp Snapshot and FlexClone® technology, along with NetApp deduplication, can support the rapid deployment of tens, hundreds, and thousands of VMs in minutes while reducing the total storage required to support such a deployment by 50% or more when compared to a baseline of traditional storage.

2.1 Provisioning Concepts

NetApp Snapshot and FlexClone

The traditional VM provisioning process involves tedious and time-consuming tasks such as provisioning storage, installing the operating system (OS) environment, patching it up with required service packs and applications, and rolling it out to the end user. NetApp storage, with its Snapshot and FlexClone technologies, facilitates an instantaneous zero-space-consuming writable copy of the flexible volumes. These FlexClone volumes can be provisioned within a matter of minutes. They contain logical unit numbers (LUNs) with virtual hard drives (VHDs) that can be connected to and used as individual OS instances for VMs.

NetApp Deduplication for Space Savings

Creating multiple copies of the golden VM images specific to user requirements can result in higher space utilization. Deduplication technology operates at the block level, eliminating data blocks with identical content and maintaining a single copy of the dataset, thereby achieving a higher level of space savings. Deduplication can be implemented in a production environment with minimal effect on storage performance and in some environments can even increase storage performance.

2.2 Virtual Machine Provisioning Process

Figure 1 shows the process flow required to provision Hyper-V VMs using NetApp cloning techniques.

Figure 1) Process flow to provision Hyper-V VMs using NetApp cloning techniques.

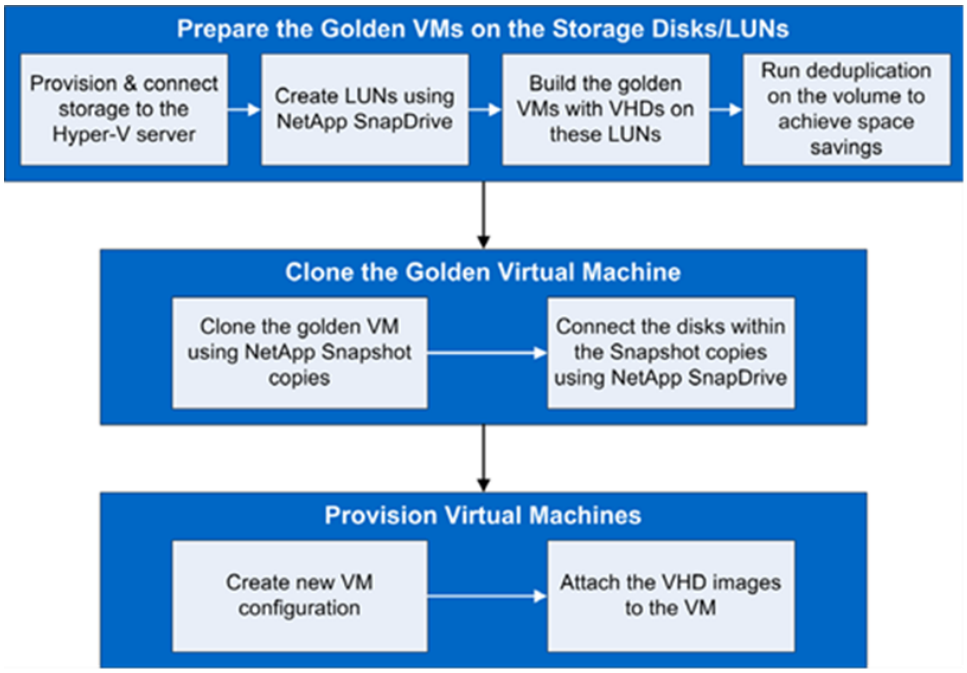

Preparing the Golden Virtual Machine

In cases in which the infrastructure requires multiple copies of the OS instance, it is repetitive and time consuming to perform installation. Server virtualization offers an efficient method to reduce the effort needed for this task and allows the attachment of an existing image as an OS disk. This facility can be used to perform a one-time installation of the child OS and designate it as a golden image. The golden image can then be updated at any time with required service packs and applications as needed. This allows administrators to provision desktops and servers to users in a matter of minutes.

To provision storage for the Hyper-V server, follow these steps:

1. Create the aggregate.

   Follow the NetApp best practice recommendations for settings for new aggregates that are described in the “Aggregates” section of NetApp Technical Report 3702, NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V.

2. Create the golden volume.

   Create a new volume within the aggregate created in the previous step and follow the NetApp best practice recommendations for settings for new flexible volumes that are described in the “Flexible Volumes” section of NetApp Technical Report 3702, NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V.

To connect storage to the Hyper-V Server:

1. Connect to the NetApp storage using NetApp SnapDrive® technology.

   NetApp SnapDrive for Windows® (SDW) can be used on the Windows Server 2008 R2 server to manage the NetApp storage system.

   a. **For details on best practices for configuration and use of NetApp SDW**, see the “NetApp SnapDrive for Windows” section of NetApp Technical Report 3702, NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V.

b. **For details on the installation of NetApp SDW**, refer to the “NetApp SnapDrive for Windows” section of NetApp Technical Report 3701, Microsoft Virtualization: Solution and Implementation Guide and follow the instructions in the “Installation” section.

If using iSCSI connectivity between the Hyper-V server and NetApp storage, establish an iSCSI session before completing the next steps.

2. Create a LUN using NetApp SnapDrive in the golden volume.

   Create a new LUN within the flexible volume created in the previous step. Then follow the NetApp best practice recommendations for settings for new LUNs that are described in the “LUNs” section in NetApp Technical Report 3702, NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V, specifically the subsection titled "Process to Create a LUN."

To build the golden VM:

3. Create a VM using Hyper-V Manager or Microsoft System Center Virtual Machine Manager (SCVMM).

   For details, see the “Virtual Machine Provisioning” section of NetApp Technical Report 3701, Microsoft Virtualization: Solution and Implementation Guide and follow the instructions in “Provisioning Using Hyper-V Manager or Provisioning Using SCVMM 2008.”

   Create a VM with a fixed VHD, of an appropriate size for the expected OS, on the LUN created in the previous step.

4. Install the child OS.

   For details, see the “Install Operating System” section of NetApp Technical Report 3701, Microsoft Virtualization: Solution and Implementation Guide and follow the instructions.

5. Once installation is complete, NetApp recommends installing the Hyper-V Integration Services, also known as integration components (ICs).

   The ICs are installed for the purpose of time synchronization, heartbeat, shutdown, key/value pair exchange, and Volume Shadow Copy Service (VSS). For details, see the “Install Hyper-V Integrated Services” section of NetApp Technical Report 3701, Microsoft Virtualization: Solution and Implementation Guide.

6. Install all applications and necessary OS and application updates.

   The VM image is a golden copy that will be further used to provision identical VMs. To avoid repetitive application installations, NetApp recommends installing all OS-related patches and required applications such as service packs, antivirus applications, and office automation software.

7. Configure the child OS and shut down the VM.

   Use the Microsoft System Preparation (SysPrep) Tool to configure VMs provisioned from a golden image before pushing them into production. This process generates a secure ID (SID) for the OS instance so that it remains unique. Refer to Microsoft KB 302577 for detailed instructions on using it.

8. Enable NetApp deduplication on the volume.

   Multiple LUNs can be created in a single FlexVol® volume, and copies of VHDs can be stored on these LUNs, which would be attached to the Hyper-V server as physical disks. Each Hyper-V VM might have the same or a different set of applications installed within the OS environment as needed. Space savings can be achieved with the NetApp deduplication capabilities. For details, follow the instructions in the “Enabling NetApp Deduplication” section of NetApp Technical Report 3702, NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V.

   Before enabling deduplication on the flexible volume, verify that best practices have been followed to disable space reservation on all LUNs within the flexible volume by unselecting Space Reserved in the LUN properties. For details, see the “LUNs” section in NetApp Technical Report 3702, NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V, specifically the subsection titled “Process to Create a LUN.”

To clone the golden VM:

9. Clone the Golden VM using NetApp Snapshot technology.

   NetApp FlexVol volumes can be cloned with zero space consumed by creating NetApp Snapshot copies. An individual connection can be established to the LUNs existing on these clones to attach them as separate physical disks to the Hyper-V server.

   Within NetApp SnapDrive, select the physical disk on which the golden VM image resides and create a NetApp Snapshot copy of this physical disk.

10. Connect the disks in the NetApp Snapshot copy using NetApp SnapDrive.

    After the NetApp Snapshot copy has been successfully created, SnapDrive can be used to connect to the individual LUNs within the Snapshot copy as individual disks. FlexClone volumes will be created from this Snapshot copy. For details, see the "FlexClone Volume Creation" section of NetApp Technical Report 3702, NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V.

To provision the VMs:

1. Create a new VM.

   The Hyper-V Server Manager or SCVMM can be used to create the desired number of VMs. For details, see the "Virtual Machine Provisioning" section of NetApp Technical Report 3701, Microsoft Virtualization: Solution and Implementation Guide and follow the instructions in "Provisioning Using Hyper-V Manager or Provisioning Using SCVMM 2008."

   When creating the new VMs, use the option `Attach a Virtual Disk Later` so that blank VM configurations will be ready for the cloned VHD.

2. Attach the golden VHD images existing on the FlexClone volume.

   Within the settings for each VM (right-click the VM), under IDE Controller 0, add a hard disk and browse to the location of the VHD stored on the NetApp FlexClone volume. NetApp recommends renaming the VHD to match the VM name or disk (for example, `[VMname]_Vol0`) before connecting it to the VM. The VM can then be powered on as usual and given a unique configuration (host name, IP address, and so on).

3 SnapManager 1.1 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 have seen their backup window grow.

NetApp SMHV addresses the resource utilization issue typically found within 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.

3.1 Leveraging NetApp Data ONTAP for Hyper-V Backup, Restore, and Disaster Recovery

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 chapters detail the best practices for deploying and using SnapManager 1.1 for Hyper-V.

3.2 Purpose and Scope

The purpose of the following chapters is to provide best practices for deploying SMHV to back up and recover Hyper-V VMs that reside on storage systems based on clustered Data ONTAP 8.1.1. They describe the key features and best practices to effectively manage the complete backup lifecycle for Hyper-V VMs. For detailed instructions on installation and configuration, refer to the SnapManager 1.1 for Hyper-V Installation and Administration Guide.

3.3 Intended Audience

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

- Data ONTAP 8.1.1 System Administration Guide
- SnapManager 1.1 for Hyper-V Installation and Administration Guide
- SnapDrive 6.4.1 for Windows Installation and Administration Guide

3.4 Technical Details

SMHV provides 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
- Allows an administrator 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 mirroring 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 crash-consistent backup and restore of virtual machines in SMHV 1.1

4 SMHV Planning

Microsoft Windows Server 2008 Server R2 with the Hyper-V role enabled offers various storage infrastructure configurations and provisioning methods. For more details, refer to NetApp Technical Report 3702, NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V.

4.1 Storage Considerations

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

5 SMHV Simplified Backup and Recovery

5.1 Prerequisites

SnapManager 1.1 for Hyper-V requires SDW 6.4 to be installed as a prerequisite if the VMs are hosted in storage systems with clustered Data ONTAP 8.1.1. If the VMs are hosted on 7-Mode systems, SnapDrive 6.2 or later is required. Refer to NetApp Technical Report 3702, NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V for more details on SMHV in 7-Mode environments.

SnapDrive manages LUNs on a storage system, making these LUNs available as local disks on Windows Hyper-V hosts. This allows Windows hosts to interact with the LUNs just as if they belonged to a directly attached redundant array of independent disks (RAID).

Note: SDW is required on Hyper-V parent hosts but not on client hosts. For WFC configurations, SDW and SMHV must be installed on each node of the cluster.

Note: SMHV 1.1 supports crash-consistent backup and restore of virtual machines. This has SnapDrive 6.4.1 for Windows as a prerequisite.

5.2 Terminology

Terminology	Description
Datasets	A dataset is a grouping 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 groupings (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).
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]). Policy includes other capabilities that make it possible to run scripts before and after the backup.
Backup and Recovery	SMHV provides local backup and recovery capability with the option of replicating backups to a remote storage system using SnapMirror relationships. Backups are performed on the whole dataset, which is a logical collection of VMs, with the option of updating the SnapMirror relationship as part of the backup on a per-job basis. Similarly, restores can be performed at an individual VM level.
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 specifying the retention period, allowing administrators the flexibility to meet varying service-level agreements (SLAs) within their environment.
Alert Notification	Alert notifications are created on a per-scheduled-backup-job basis and are sent by e-mail to administrator-defined accounts. Alert notification can be configured to e-mail

Terminology	Description
	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/Restore

These backups are taken in coordination with Volume Shadow copy Service (VSS) to make sure that the applications running in the VM are quiesced before taking a Snapshot copy. Such a backup guarantees the integrity of application data, and hence can be safely used to restore the VM and the applications running in the VM to a consistent state.

Crash-Consistent Backup

A backup in which the state of data is equivalent to what would be found following a catastrophic failure that abruptly shuts down the system. The data in the backup will be the same as it would be after a system failure or power outage. This type of backup is much quicker. A restore from such a backup would be equivalent to a reboot following an abrupt shutdown.

Note: Crash-consistent backup and restore is supported from SMHV 1.1 onward and will require SnapDrive for Windows 6.4.1 to be installed on the host system.

5.3 Port Usage

Best Practice

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

- 808: SMHV and SDW default port
- 4094: If SDW is configured to use 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.

5.4 Architecture

Figure 2 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 2) SMHV architecture.

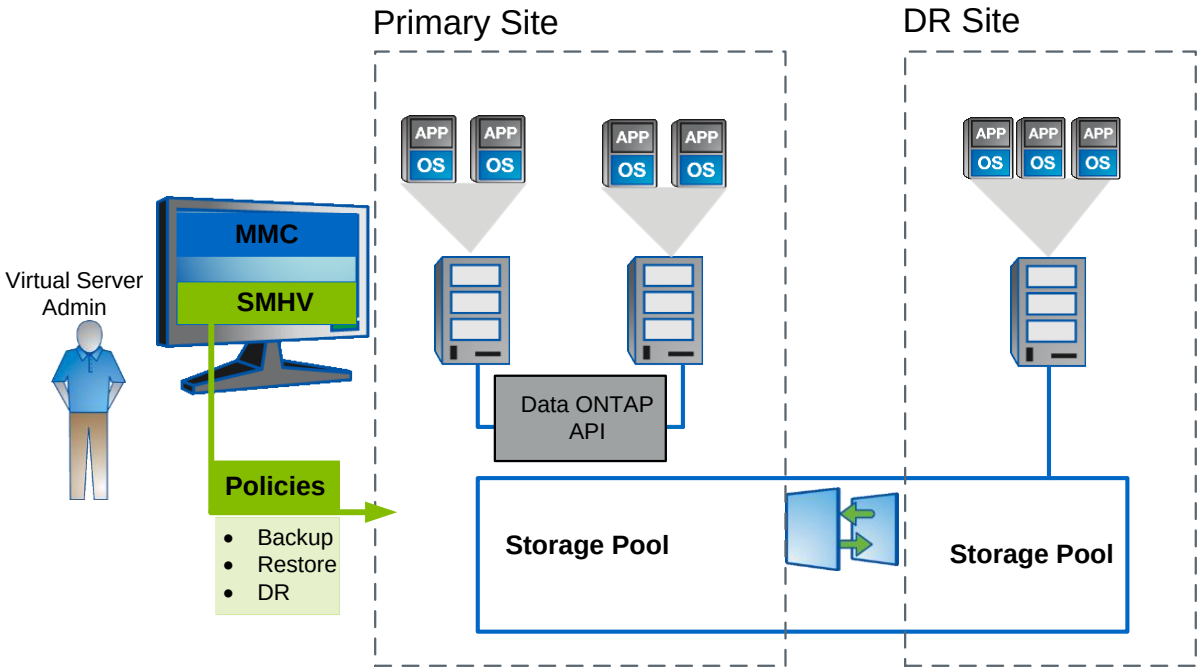

Components

License Requirements

An SMHV license is required on the Windows host system. There is a choice of either host-based or storage-based licensing.

- **Host-based licensing** requires that a license key be provided during installation. The license key can be changed after installation by clicking License settings in the SMHV Welcome window.
- **Storage-based licensing** requires that the SMHV license be added to all storage systems.

Systems with NetApp Clustered Data ONTAP 8.1.1

SMHV requires a different version of software in order to manage VMs hosted on LUNs located on storage systems with clustered Data ONTAP 8.1.1 supports both host-based and storage-based licensing.

Table 1 lists the licensed and supported versions of Data ONTAP 8.1.1.

Table 1) Licensing and supported versions of clustered Data ONTAP 8.1.1.

7-Mode	Clustered Data ONTAP
Data ONTAP 8.0 or earlier	Data ONTAP 8.1.1 or later
Device-specific module (DSM) 3.4 or earlier	DSM 3.5
SnapDrive 6.3 for Windows or earlier	SnapDrive 6.4 for Windows

7-Mode	Clustered Data ONTAP
Windows Host Utility Kit 5.0	No longer required for Data ONTAP DSM. The DSM 3.5 Installer includes the MBRAlign tool and the `LinuxGuestConfig.iso` in the installation package. If MSDSM is installed, then Windows Host Utilities 6.0 must be installed.
SnapManager for Hyper-V 1.1 P1	SnapManager for Hyper-V 1.1 P1 (no change)

Figure 3 shows how SMHV can be deployed to manage virtual entities in an environment with clustered Data ONTAP 8.1.1.

Figure 3) SMHV deployed to manage virtual entities in a clustered Data ONTAP environment.

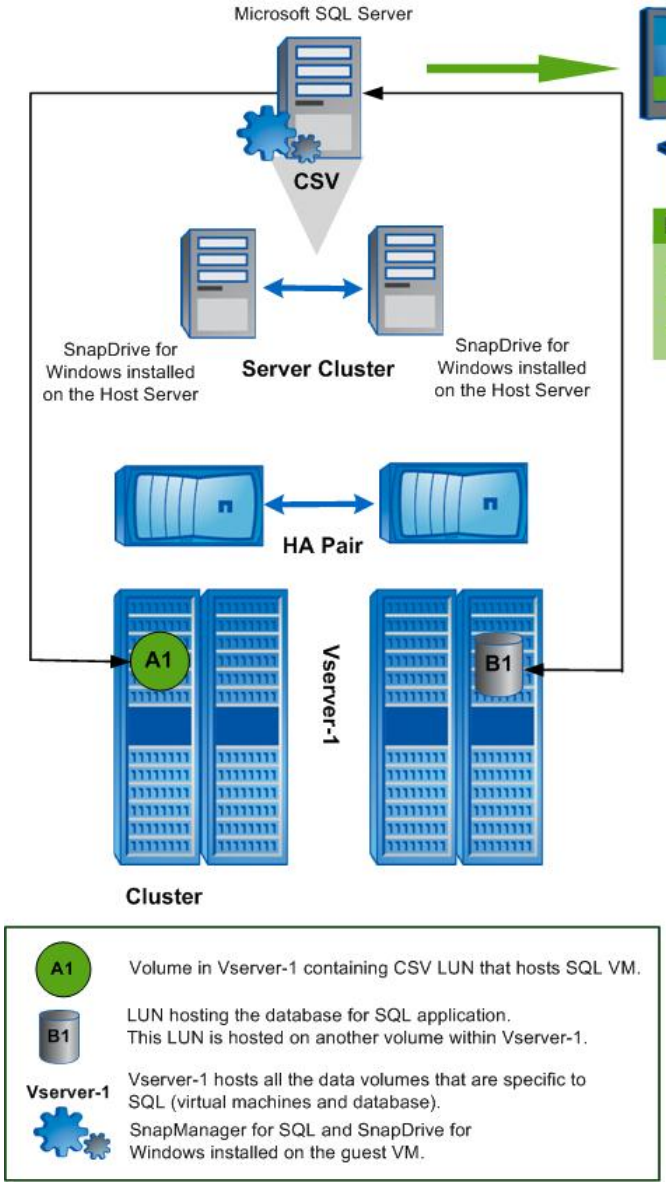

SMHV-Supported Configurations

SMHV must run on Windows Server® 2008 R2 x64.

Platform Support

- Windows Server 2008 R2 x64 Standard, Data Center, Enterprise editions (full and core installation)
- Hyper-V Server 2008 R2 x64

Remote Management Platform Support

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

VM Support

- Windows Server 2008 R2 x64 (all editions, full and core)
- Windows Server 2008 x64 Standard, Enterprise (full and core)
- Windows Server 2008 x64 Standard, Enterprise with SP2 (full and core)
- Windows Server 2003 x64 and x86 with SP2 and later
- Windows Vista
- Windows XP
- SuSE Linux® (SLES10 SP 1 and SP2) x86 and x64
- RHEL 5.3, RHEL 5.4, and RHEL 5.5 (Microsoft Hyper-V Integration component version 2.1 must be installed)

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

SMHV SnapInfo Settings

SMHV SnapInfo folder 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.

Note: The SnapInfo path must reside on a Data ONTAP LUN. For managing dedicated VMs, the SnapInfo location must be a dedicated Data ONTAP LUN. For managing shared VMs, the SnapInfo location must be a shared Data ONTAP LUN.

The SnapInfo path must not reside on a CSV.

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

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

Best Practice
For clustered Data ONTAP 8.1.1, NetApp recommends having the SnapInfo LUN in a separate volume within the Vserver and not as part of the other data volumes. For example, if the SMHV is protecting SQL VMs that are hosted in a CSV LUN in a volume, then the user must make sure that the SnapInfo LUN is not part of this volume of the Vserver. This simplifies VM restoration and DR.

SMHV Report Settings

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

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

SMHV Event Notifications

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

6 SMHV Process Flow

6.1 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.

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 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 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 a previously exported setting. If this operation is performed in a clustered environment, the settings on all nodes in the cluster must be imported so that all host and dataset configurations are the same.

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

6.2 The Backup Process and Implications

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 physically reside.

In addition to the Snapshot copy stored locally, SMHV also provides an option to update an existing SnapMirror relationship at the completion of a backup. This can be selected on a per-backup-job basis as required by the administrator. 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.

Figure 4 provides a high-level overview of the typical SMHV architecture on the primary site storage, where the backup process takes place.

Figure 4) Hyper-V infrastructure and associated storage during an application-consistent SMHV backup.

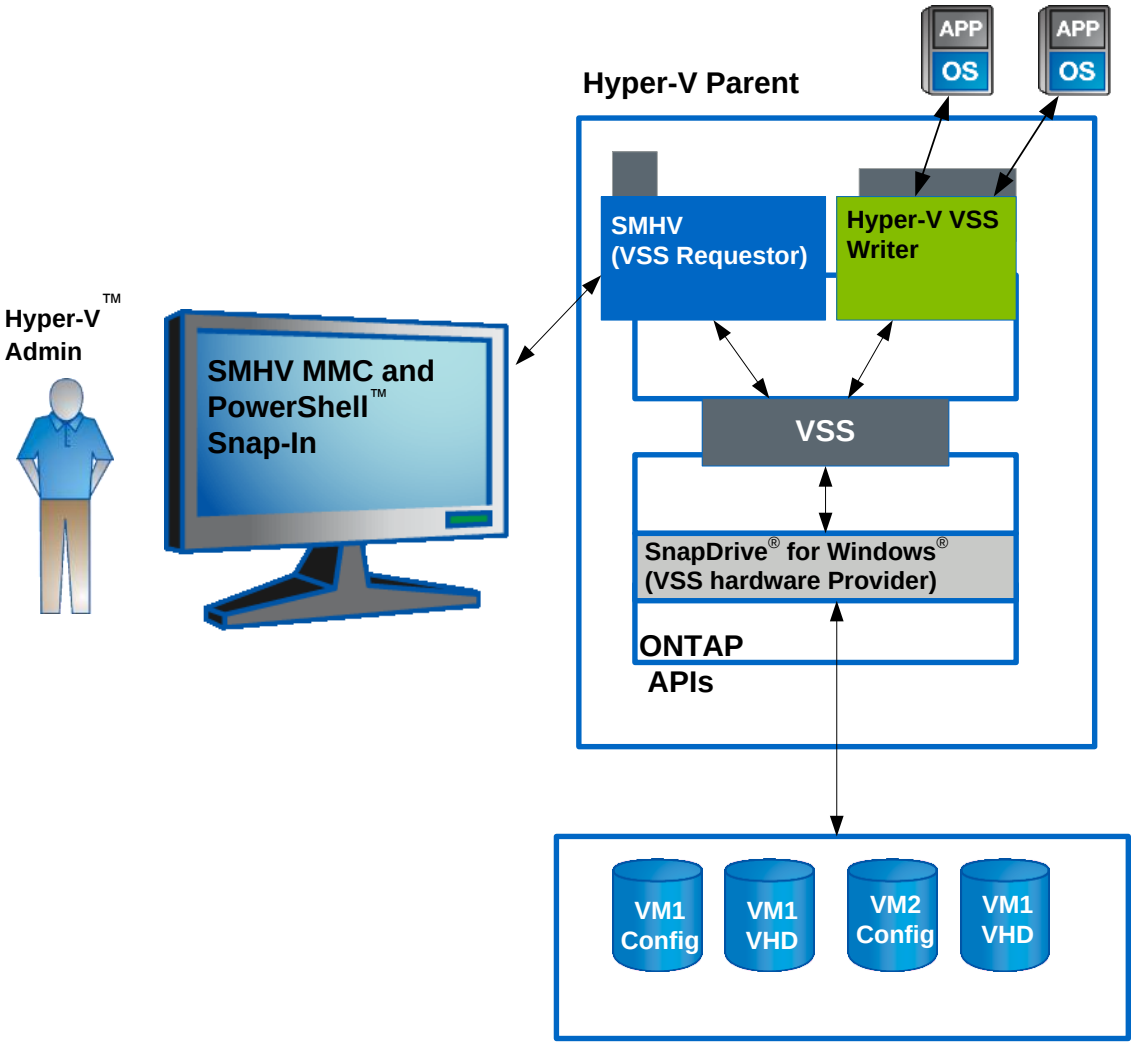

The following steps describe the flow of the backup process.

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 then implements a VSS requestor component to coordinate the backup process and create a consistent Snapshot copy in Data ONTAP using VSS hardware provider for Data ONTAP LUNs.
4. The VSS framework asks the hardware provider to mount the LUNs from the Snapshot copy.
5. The Hyper-V writer recovers data on the LUNs and brings it to the state of the software Snapshot copy that was created in step 2.
6. The VSS provider creates a second Snapshot copy of the LUNs and then dismounts them from the Snapshot copy.
7. At the completion of the local backup, SMHV updates an existing SnapMirror relationship on the volume if the SnapMirror option was selected. SnapMirror is discussed in further detail in section 8, "SMHV Disaster Recovery."

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 using the option Copy the contents of the specified physical disk in the Configure Disk page in the new VHD creation wizard in Microsoft Hyper-V Manager.

Note: SMHV does not support the backup and restore of VMs running on SAN boot LUNs. This is a limitation of SDW.

Workflow for Crash-consistent backups:

1. User chooses crash-consistent backup option in the backup dataset wizard.
2. 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 take a Snapshot copy of these LUNs. Only one Snapshot copy will be taken for each LUN irrespective of the number of VMs running on it.
4. Backup will be registered with backup type as 'Crash-consistent.'
5. Upon completion of the local backup, SMHV updates an existing SnapMirror relationship on the volume if the SnapMirror option was selected.

Note: While performing a crash-consistent backup or restore, SMHV 1.1 does not leverage VSS. VSS is used only to get VM-related metadata from the Hyper-V writer. The default backup type will be application-consistent backup.

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 using Hyper-V Manager before creating a backup using SMHV. If this is not done, the backup could fail.

Best Practice

For clustered Data ONTAP 8.1.1, all VMs related to an application can be isolated to a single Vserver. A single dataset can be created to back up and restore these VMs. This simplifies the backup, restore, and DR processes.

6.3 Scheduled Backups and Retention Policies

SMHV allows administrators to schedule a dataset backup at a particular time. SMHV uses the Windows Tasks Scheduler for creating or modifying scheduling policies. The 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 so as 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

Since the SnapManager suite of products (SnapManager for SQL, SnapManager for SharePoint, SnapManager for Exchange, 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 will reduce the performance overhead on the cluster server (C-server).

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 not only to meet specific SLAs, but also to maintain a supported number of NetApp Snapshot copies on the underlying volumes. For SMHV, one backup creates two Snapshot copies on the storage systems for data consistency (refer to KB ID: 2010607). For example, setting a retention policy of 30 backups on an hourly 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 will be reached in 5 days, and backups will begin to 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 they are manually deleted by the administrator. 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.

After creating a dataset backup, SMHV creates a Snapshot copy of the SnapInfo LUN. SnapInfo Snapshot copies are not deleted if the backup is deleted. SnapInfo Snapshot copies have a different retention policy. By default, SMHV retains 30 SnapInfo LUN Snapshot copies and deletes the older ones when the SnapInfo Snapshot copy count exceeds 30. In configuring the number of SnapInfo Snapshot copies to be retained for each Hyper-V host, use the following registry keys:

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

6.4 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 this 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 check box allows SMHV to back up the VM using the saved state. If this option is checked, 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 or 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: 'Allow saved state policy' option is not applicable for crash-consistent backups. This is because the VM is being backed up irrespective of the state.

6.5 Backup Scripts

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 LUN at the end of the dataset backup.

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

Note: The `$SnapInfoSnapshot` variable is supported for dedicated VMs only.

6.6 Quick/Live Migration Implications

Best Practice
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.

6.7 Restore Process

SMHV can restore a VM 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 SDW. The associated files of a VM, including the configuration file, Snapshot copies, and any VHDs, can be spread across multiple Data ONTAP LUNs. A LUN can contain files belonging to multiple VMs.

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

With these differences in restore types aside, the following process flow is used by SMHV 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:

- The VM to be restored is not [currently running] on the host.
- 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.
- The Snapshot location of the VM has changed.
- The names of VHD files or their file system paths or NetApp storage system LUN paths have changed.

In all of these warning scenarios, the VM can be restored, but 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 same, the restored VM might have additional or fewer VHDs. If that is the case, NetApp recommends that the cluster configuration of the VM and its dependencies be manually updated.

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

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 addition, RestoreToAlternateHost switch in Restore-Backup cmdlet cannot be used when the backup being restored is a crash-consistent backup.

6.8 Mounting a Backup

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

Perform these steps to mount a backup:

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

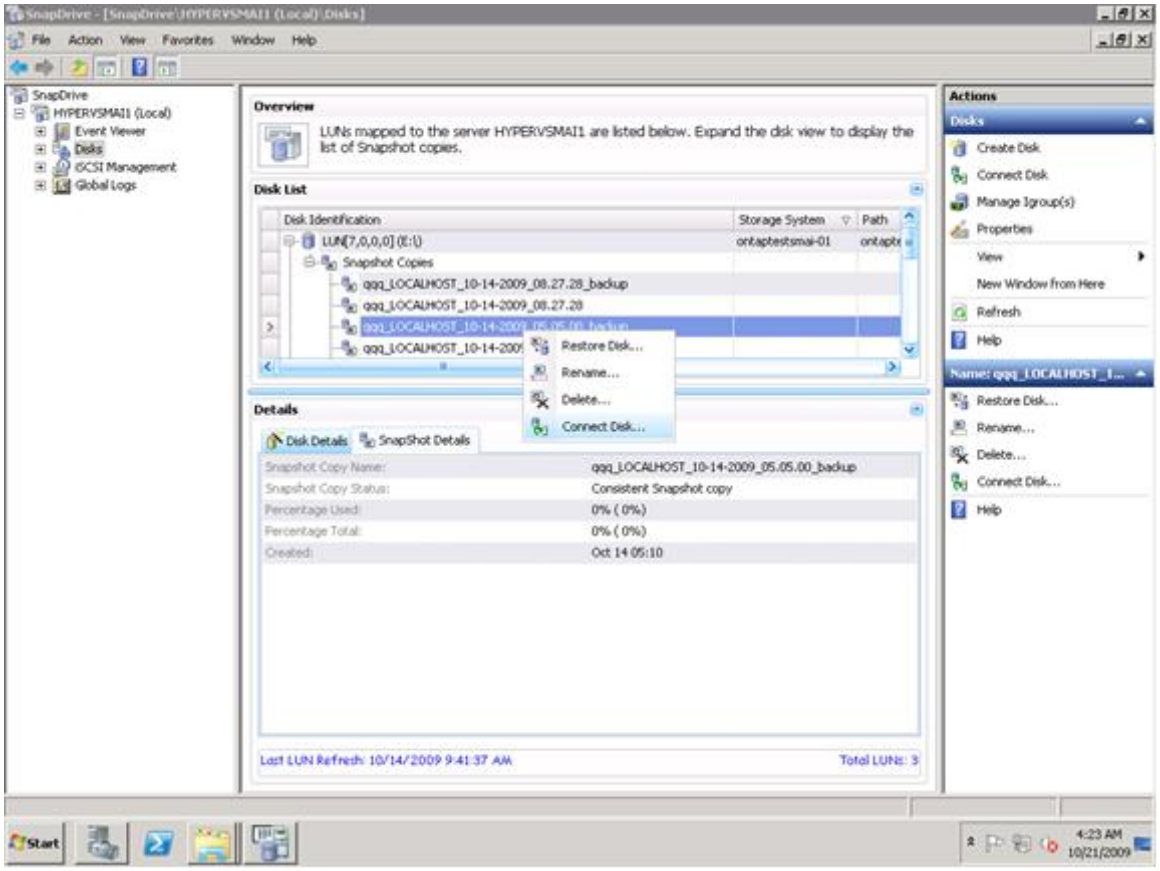

2. Right-click the Snapshot copy (the one with the _backup suffix) and select the Connect Disk option.
3. Click Next.
4. If the LUN is a dedicated disk, proceed to the next step; otherwise, if the LUN is a Windows cluster resource, perform the following steps in the Specify Microsoft Cluster Services Group panel. In the Specify Microsoft Cluster Services Group panel, 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 option Create a new cluster group.

     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.

   - To CSVs, select the Add option.
5. 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\).
6. In the Select Initiators panel, choose an initiator for the LUN.
7. In the Select Initiator Group Management panel, specify either automatic or manual igroup management.
8. In the Completing the Connect Disk Wizard panel, 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.

9. 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 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 recently mounted in the preceding step using the Disk Manager snap-in. Select the disk and choose the offline menu item. This offlines the disk mounted from the VHD.
3. Attach the offlined disk to the VM by selecting the Physical hard disk button. Then select the offlined disk from the drop-down list under the Physical hard disk button on the Settings for Exchange_Production screen. This presents a new drive inside the VM (backed by VHD in the parent).

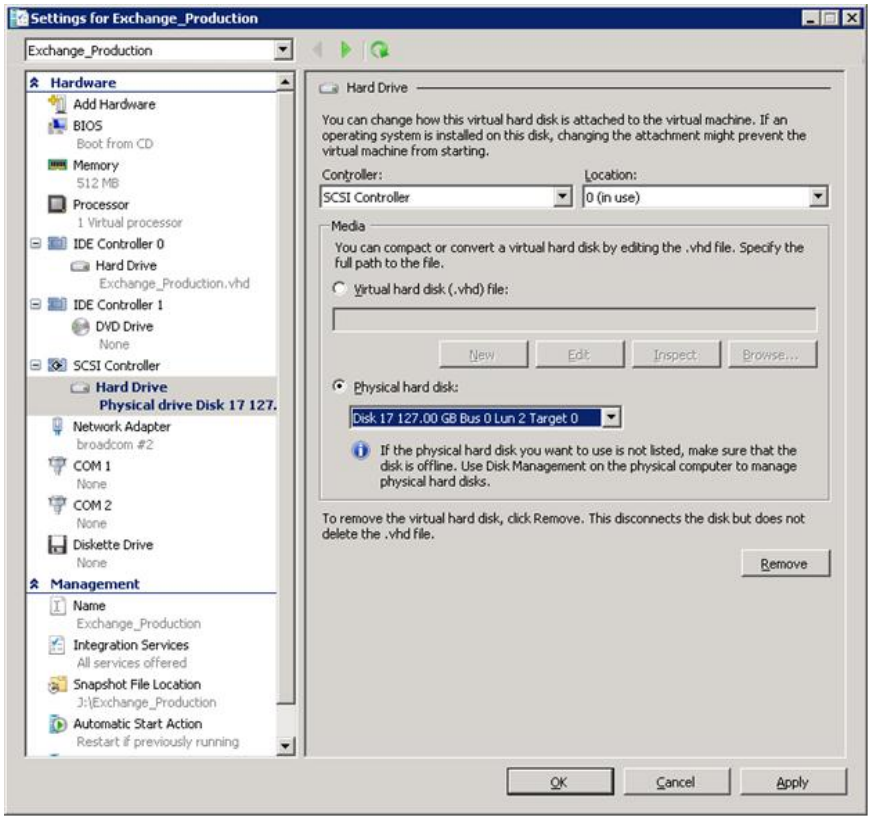

4. Log into the VM and select the newly mounted drive to see the contents of the disk backed by the VHD attached.
5. When verification is complete, detach the disk from the VM, using the VM settings screen, and click Remove. Use SDW to unmount the disk backed by the Snapshot copy, using the SnapDrive

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

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

7 SMHV 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 the SMHV, as a host-end application, offers services provided that the storage is continuously available. Section 7.1 provides detailed description of the available tools that facilitate storage availability.

7.1 Multipath High Availability with Active-Active NetApp Controllers

The NetApp active-active controllers offer easy, automatic, and transparent failover capabilities to deliver an 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 storage 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.

7.2 Data ONTAP DSM for Windows MPIO

Microsoft MPIO is a protocol-independent feature that supports multiple data paths to a storage device with iSCSI, FC, or SAS. Providing multiple paths that can handle failover increases the availability from a host to the storage system. Windows Server 2008 R2 x64 servers include support for Microsoft MPIO.

NetApp Data ONTAP DSMs for Windows MPIO help NetApp storage systems to integrate with Microsoft MPIO on Windows Server 2008 R2 servers and provide HA to applications using path-failover methods. The DSM determines all the paths pointing to the same LUN so that MPIO can group them into the virtual disk that the Windows Server 2008 Hyper-V server will mount. It is also responsible for communicating with MPIO to identify the path to which to route I/O. This is especially important in the event of a failover. There can be multiple active paths and multiple passive paths. If all of the active paths fail, the DSM automatically switches to the passive paths, maintaining the host's access to its storage.

Best Practices

- For a highly available connection to the storage system, NetApp requires installing the supported version of multipathing software such as the Data ONTAP DSM for Windows MPIO.
- For Windows Server 2008 R2 servers, NetApp recommends Data ONTAP DSM 3.2R1 or higher.
- For the currently supported multipathing software versions and related requirements, see the NetApp Interoperability Matrix Tool.
- For clustered Data ONTAP 8.1.1, refer to the DSM versions requirement in Table 3.

8 SMHV Disaster Recovery

The DR functionality can be used after updating the SMHV software with a patch available at the NetApp Support (formerly NOW®) site.

After upgrading to SMHV 1.0 P1, the user can perform failover and failback of Hyper-V VMs using Windows PowerShell™ cmdlets in the SMHV PowerShell option. The Windows PowerShell cmdlet `restore-backup` must be used along with the switch `–RestoreToAlternateHost` and the server name to use this feature.

For example:

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

New Cmdlet: Get-VMsFromBackup

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

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

For example:

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

Basic DR Scenario

Components

- 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.

Steps to Fail Over VMs to Secondary

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. The LUN type and mount point must be the same as that of the primary. 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 taken 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.
6. Use the Restore-backup cmdlet with VM GUID (from step 5) and backup from (step 6). 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.

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

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.

Steps to Fail Back VMs to Primary

Following are the steps to fail back VMs to primary:

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 temporarily down, 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 (UI) 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 connecting 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: Taking 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 graphical user interface (GUI). Most users will probably take the SMHV backup in lieu of manually performing a mirror update because it can be scripted, whereas the mirror update is a tedious job and prone to user error (such as forgetting to update a LUN).

5. Connect to all LUNs on the primary (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 can be released.

6. Restore the SnapInfo LUN from its last Snapshot copy taken by SMHV.
7. Add the primary host or cluster in SMHV MMC and configure it with the SnapInfo path.
8. Use the Get-VMsFromBackup cmdlet to get a list of VMs present in backup metadata.
9. Use the Get-Backup cmdlet to get the backups for each VM.
10. Use the Restore-backup cmdlet with VM GUID (from step 6) and backup from (step 7). Use -RestoreToAlternateHost switch and specify the primary host or cluster name as -server parameter. If the primary 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 using failover cluster UI/Windows PowerShell cmdlets.

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 disk and the disks containing VMs using 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 using SnapDrive.
2. Resync the SnapMirror relationship from the primary to the secondary.

Updating SnapMirror in SMHV for Virtual Machines Residing in Systems with Clustered Data ONTAP 8.1.1

From an administrative perspective, there is no change in the way SMHV must be configured for updating SnapMirror relationships. In the case of clustered Data ONTAP systems, to update SnapMirror in SnapManager for Hyper-V, both the cluster IP (C-server) credentials and the Vserver credentials of the source volume and destination volume in SDW 6.4 must be added.

Note: SMHV and SnapDrive support both replication across clusters and replication within clusters. In the case of intercluster replication, at least one intercluster logical interface (intercluster LIF) is required per node. The intercluster LIF can be assigned to a data port or a dedicated intercluster port.

For more details on SnapDrive configuration for data protection, refer to "SnapDrive 6.4 for Windows Best Practices Guide" and to "Installation and Administration Guide for Data Protection in Cluster-Mode."

9 SMHV Application Consistency

Microsoft's VSS was written 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 within 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 are resumed. 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, or any other VSS-aware application is running on VHDs in the VM. The applications that exist in the VM restore 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, 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 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 details on VM backup, refer to Planning for Backup on the Microsoft TechNet site.

10 Crash-Consistent Backup and Restore

Backups taken using SMHV 1.1 can be either application-consistent or crash-consistent. Application-consistent backups are taken in coordination with Volume Shadow Copy Service (VSS) to make sure that the applications running in the VM are quiesced before taking the Snapshot copy. Such a backup assures the integrity of application data; hence, can be safely used to restore the VM and the applications running in the VM to a consistent state.

Though 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. Since the application writer in the VM and the Hyper-V writer in the parent OS are involved, the backup process. Failure in any of the components will fail the backup.
- Hyper-V writer uses the auto-recovery process to make the VMs consistent. Auto-recovery results in the creation of two Snapshot copies on 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 will be desirable to have some way of taking "quick" Hyper-V VM backups. Crash-consistent backup is designed to provide this ability of taking quick backups.

A crash-consistent backup of a VM will neither use VSS to quiesce data, nor will it result in auto-recovery. This backup will simply take 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 will be the same as it would be after a system failure or power outage. All the SMHV functions such as scheduling, restore, script execution, SnapMirror updates, backup retention, and so on will be supported for crash-consistent backups as well.

The following screenshot shows the backup dataset wizard showing backup types: “Application-consistent and Crash-consistent.”

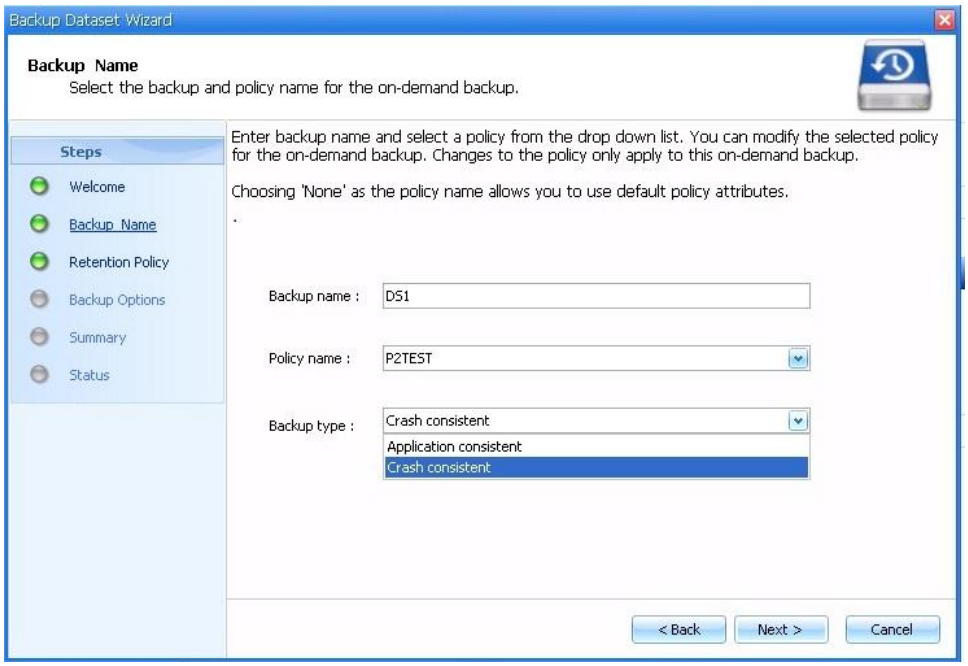

Note: Saved state backup policy is not applicable for crash-consistent backup and restore. This is because crash-consistent backups do not involve the Hyper-V VSS writer.

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

Best Practice
The crash-consistent backup feature is not a replacement for application-consistent backups. It enables user to have frequent recovery points. Therefore, user can have frequent crash-consistent backups and fewer application-consistent backups.

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

11 Windows Server 2012 Support

Windows Server 2012 supports SnapManager for Hyper-V 1.2 onward. SnapDrive 6.5 for Windows is prerequisite software that must be installed on Windows Server 2012 to use SMHV. With SnapManager for Hyper-V 1.2, backup and restore of virtual machines will be supported only in SAN environments for Windows Server 2012.

11.1 Prerequisites

The prerequisites for SnapManager for Hyper-V 1.2 are as follows:

- **SnapDrive 6.5 for Windows.**
- **Microsoft Device-Specific Module (MSDSM) (for multipathing).** SnapDrive for Windows operating in Windows Server 2012 does not support Data ONTAP DSM. For multipath I/O (MPIO) operations, use MSDSM.
- **Windows Host Utilities Kit 6.0.1 (mandatory).** It is mandatory to install Windows Host Utilities kit 6.0.1 on the host and the guest VM. After installation, Windows Server 2012 space reclamation is disabled. Space reclamation for NetApp storage LUNs should be performed using SnapDrive 6.5 for Windows.
- **.Net 3.5.1**. Windows Server 2012 has .Net 4.0 as well as .Net 3.5. The user is required to install .Net 3.5 for SnapDrive 6.5 for Windows.

11.2 Feature Overview

SnapManager for Hyper-V 1.2 with SnapDrive 6.5 for Windows will support all major SAN-based features in Windows Server 2012.

Here is an overview of all the features and best practices to be followed.

CSV 2.0 Support (CSVFS)

Note: In Windows Server 2012, CSVs have undergone significant changes with respect to security, performance, and file system availability for additional cluster workloads. A new clustered file system has been introduced, and this functions as a layer of abstraction above the NTFS file system for the storage volume. As a result, simultaneous reads/writes can be performed on the CSV LUN from different nodes. For more details on CSV 2.0, refer to

http://technet.microsoft.com/en-us/library/jj612868.aspx. A CSV 2.0 volume will have two volume GUIDs.

- **NTFS volume GUID.** When a disk is created and partitioned with NTFS and before it is added to the CSV.
- **CSV volume GUID.** When a disk is added to the CSVs.

Best Practice
NetApp recommends, in SDW, creating a CSV from the node that owns the available cluster storage group. Use the "`CLUSTER GROUP`" command or "`Get-Cluster Group`" cmdlet to identify the node that owns "Available Storage" group before creating a CSV disk.

SnapManager for Hyper-V 1.2 supports virtual machines hosted on CSVFS volume type. The new CSVFS volume type has introduced a new CSV writer and CSV shadow copy provider. This has facilitated achieving distributed application-consistent backups. Section 11.3, “Asymmetric Clustering,” covers distributed application-consistent backup in detail.

11.3 Asymmetric Clustering

Asymmetric clustering is a feature with which users can create a shared disk or CSV among only a few nodes in a cluster.

Note: SnapManager for Hyper-V 1.2 does not support having virtual machines in such CSVs.

11.4 BitLocker Encryption

BitLocker was a data protection feature and was part of Windows 7 and Windows Server 2008 R2. This feature is now available with Windows Server 2012 with the additional functionality. The user will now be able to encrypt cluster shared SAN volumes. For more information on BitLocker configuration, refer to http://technet.microsoft.com/en-us/library/hh831713.

SnapManager for Hyper-V will support BitLocker functionality for CSVs provisioned through SnapDrive 6.5 for Windows. Virtual machines can be hosted in encrypted CSVs.

11.5 New Virtual Hard Disk Format

Windows Server 2012 has introduced a new virtual hard disk format, VHDx. Unlike the previous VHD format, this format supports up to a 64TB size. Also, the VHDx format has a 4kB logical sector size that increases performance of applications that are designed for 4kB sector sizes.

SnapDrive 6.5 for Windows supports this new format. The block allocation unit size of LUNs created by SnapDrive is 4kB. This complements the new VHDx format, and there is no scope for VM misalignment.

SnapManager for Hyper-V 1.2 will support backup, restore, and replication of virtual machines in VHDx format.

Note: SnapDrive 6.5 for Windows currently cannot create LUNs beyond 16TB, and, therefore, NetApp advises creating a VHDx for sizes less than 16TB and to use other means of provisioning additional storage (pass-through disks, guest iSCSI initiator) on the VM.

11.6 Hyper-V Virtual Machine Live Migration

In Windows Server 2012, users can perform concurrent live migration of multiple VMs from one node to another.

Best Practice
It is best to avoid SMHV-related operations within the virtual machine during live migration.

11.7 Hyper-V VM Storage Live Migration

This feature in Windows Server 2012 enables migrating virtual machine–related files to a different storage location without the VM having to undergo downtime. It is no longer necessary to take the virtual machine state offline when migrating to a different storage system.

Note: After migrating the virtual machine from one volume to another, restoring to a Snapshot copy taken in the earlier volume is not supported.

Best Practice
It is best to avoid SMHV-related operations during storage live migration. Otherwise, such operations could corrupt the virtual machine.

11.8 Windows Server 2012 Features Not Supported from SnapManager for Hyper-V 1.2 and SnapDrive 6.5 for Windows When Connected to NetApp Storage Systems Running in Clustered Data ONTAP Systems

NetApp Data ONTAP, SnapDrive 6.5 for Windows, and the NetApp SnapManager suite of products do not support the following features for Windows Server 2012:

- Hyper-V over SMB 3.0
- SMB over remote file shares
- SMB VSS for remote file shares (remote VSS)
- Virtual Fibre Channel
- Hyper-V replica
- Windows Server 2012 native thin provisioning
- Offload data transfer capability

12 SnapManager for Hyper-V 1.2 Backup Mechanism for Windows Server 2012

In Windows Server 2012, Microsoft introduced the CSV proxy file system (CSVFS). The CSVFS 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 (in Windows Server 2008 R2). For additional information on CSVFS architecture, refer to this link.

In Windows Server 2008 R2, CSV Hyper-V backup creates application-consistent backups on the 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 is taken from the Hyper-V parent using the Data ONTAP VSS hardware provider (SnapDrive for Windows). This resulted in 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 was increased.

In Windows Server 2012, CSVFS introduces “distributed application-consistent backups.” This allows backup of all the VMs in a cluster to be consistent in “one single application-consistent backup.” In order 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 back 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 auto recovery process.

Figure 5 illustrates SMHV 1.2 backup process for Windows Server 2012.

Figure 5) SMHV 1.2 backup process for Windows Server 2012.

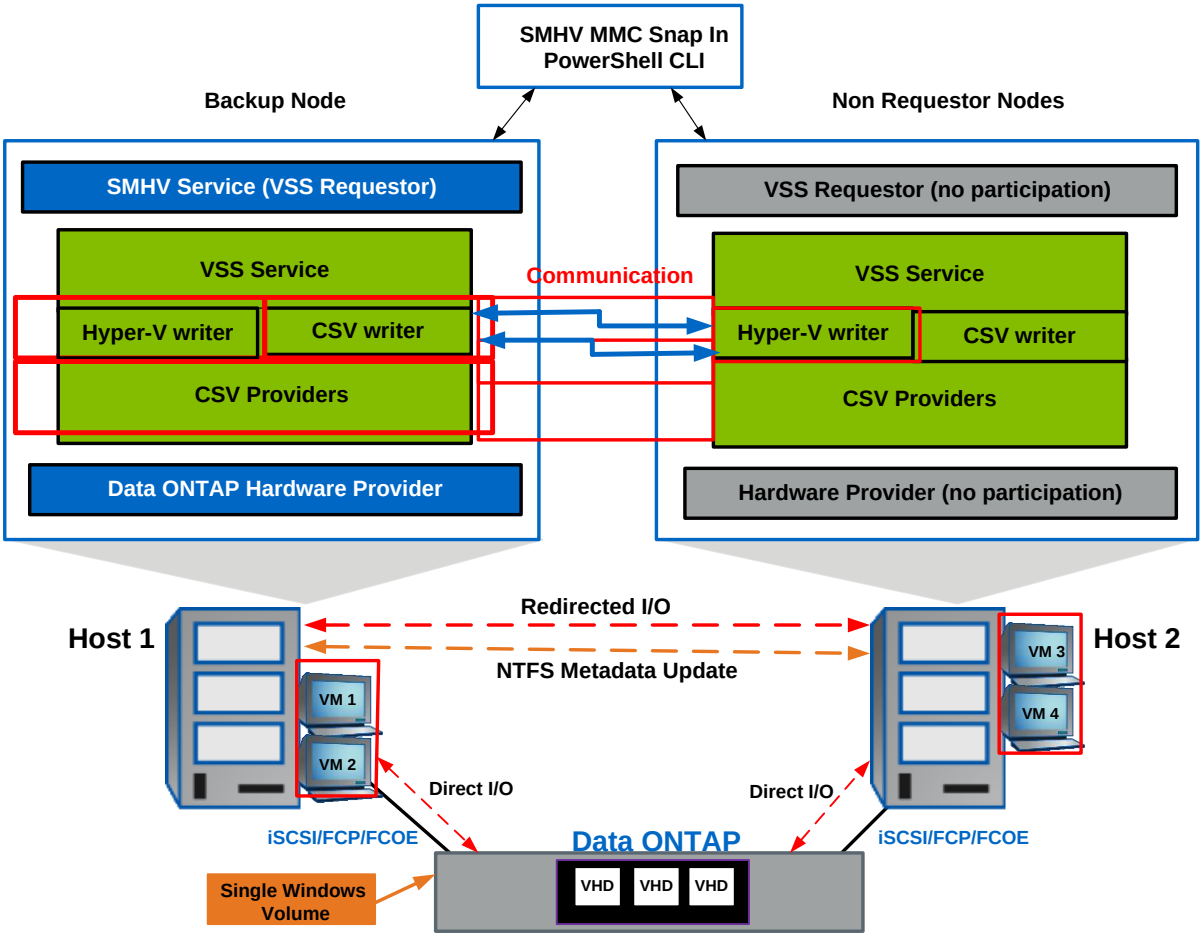

Initialization Phase

- The user initiates the backup operation from any node in the cluster 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, wherein VSS is initialized in each node of the Windows cluster, which 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

- Hyper-V writer on the coordinator node quiesces the application writers inside the VM using the integration service.
- 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 VM is consistent before starting the actual Snapshot copy of the volume.

Backup Phase

- VSS hardware provider on the coordinator node takes the backup Snapshot copy of the CSV volume.
- Hyper-V writer, by default, performs an autorecovery process on each VM owner node after the hardware Snapshot copy is created to remove any inflight 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 CSV.

Postbackup Phase

- SMHV retrieves the VSS backup metadata and backup component documents and then modifies both the metadata to make it compatible with VSS required semantics.
- SMHV saves the backup metadata to the snapinfo folder.
- VSS Snapshot GUID is renamed to 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 checked 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: It is recommended that all the VHD files belonging to a virtual machine are hosted on CSVFS LUNs only and not a mix of CSVFS and shared disks. This is because SMHV does not support such mixed-mode backups.

Best Practice

In order to achieve a successful backup and faster backup performance, it is recommended not to have more than 15 CSVFS LUNs in a single SMHV backup dataset that belong to the same NetApp storage system. In other words, virtual machines hosted on not more than 15 CSVFS LUNs belonging to the same storage system should be grouped together in a single dataset.

If we have 20 CSVFS LUNs hosted on a single NetApp storage system, it is recommended to create two datasets minimally and spread the virtual machines (CSVFS LUNs) evenly across these datasets.

To summarize, distributed application-consistent backups are faster since 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 since 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 could be stored for a volume, so considering that aspect, this enhancement would allow storing more backups for a VM.

13 Summary of SMHV Best Practices

Best Practice

When Hyper-V is deployed on shared storage, NetApp recommends configuring one VM per LUN. All VHDs related to a single VM (a VM with multiple drives) can reside on a single LUN provisioned as shared storage to a WFC. This is a best practice for Windows Server 2008 Server R2 running Hyper-V deployed on standard shared-storage volumes.

Best Practice

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

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

The default port number is 808. When SMHV is installed on a cluster, the same port number should be used across all nodes.

Best Practice

Having a SnapInfo LUN on a volume of its own is preferable.

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 reduce space consumption on the storage system.

Best Practice

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

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 to avoid exceeding 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

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 set from which to restore in case of a disaster.

Best Practice

For mission-critical VMs, NetApp recommends enabling the Allow Saved State VM Backup option.

Best Practice

SMHV cannot back up a VM that is actively undergoing migration. If a backup runs against a dataset that has VMs being actively migrated, an error is generated, and those particular VMs are not backed up. NetApp recommends that VMs be migrated only when a significant gain in performance can be achieved. This will improve not only the success rate of the backups, but the overall VM performance as well.

Best Practice

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

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 NetApp Technical Report 3450, Active-Active Controller Overview and Best Practices Guidelines.

Best Practices

- For a highly available connection to the storage system, NetApp requires installing the supported version of multipathing software such as the Data ONTAP DSM for Windows MPIO.
- For Windows Server 2008 R2 servers, NetApp recommends Data ONTAP DSM 3.2R1 or higher.
- For the currently supported multipathing software versions and related requirements, refer to the NetApp Interoperability Matrix Tool.

14 SMHV Conclusion

SnapManager 1.1 for Hyper-V provides a rich feature set that allows IT organizations to take advantage of NetApp Snapshot and SnapMirror technologies to provide fast, space-efficient disk-based backups in a

Hyper-V environment with NetApp storage while placing minimal overhead on the associated virtual infrastructure. The recommendations and examples in this report will help administrators get the most out of SMHV deployments.

Appendixes

Table 2 lists the terminology used in clustered Data ONTAP 8.1.1.

Table 2) Terminology used in clustered Data ONTAP 8.1.1.

Terminology	Description
Cluster	In clustered Data ONTAP 8.1.1, 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.
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 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 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).
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 Vserver has a separate administrator authentication domain and can be managed independently by a Vserver 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 the cluster and Vserver administrators to manage distributed data storage as a single file system.

Quick Steps to Deploy Clustered Data ONTAP Storage System

Follow these steps to deploy a storage system using clustered Data ONTAP 8.1.1:

1. Set up the cluster environment (refer to Data ONTAP 8.1.1.1 Installation and Administration Guide).
2. Create an aggregate.
3. Create a Vserver. (The following screen shows sample 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 Vserver:
   - The data LIFs, which enable Vservers to serve data to the clients (iSCSI, FCP)
   - The management LIF that will allow 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 Vserver of the secondary Vserver. Make sure the property of that volume is of the data protection (DP) type.
   b. Establish a SnapMirror relationship between the primary and the secondary by accessing the secondary system.

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

   For more information on data protection, refer to Data Protection Best Practices Guide.
8. After creating the data volumes in the storage system, create clustered LUNs of the desired size using SDW 6.4.
9. Create and host the required VMs in the LUNs.

Quick Steps to Deploy a Windows Server 2008 R2 Hyper-V Cluster Environment on NetApp Storage

Follow the steps for each of the following tasks to deploy a Windows Server 2008 R2 Hyper-V cluster environment on NetApp storage.

1. Install the NetApp storage system:
   a. Create aggregates to support the infrastructure
   b. Create volumes to support a CSV infrastructure (turn on thin provisioning)
2. Perform server operating system preparation:
   a. Install Windows OS:
   – Hyper-V role

- Failover cluster feature
- NET 3.5 feature
- MPIO feature
- All patches

b. Install Microsoft hot fixes:

- KB ID: 975921
- KB ID: 974909
- KB ID: 975354
- KB ID: 979743
- Install Windows Server 2008 R2 SP1
- KB ID: 2406705
- KB ID: 978157

c. Install NetApp software:

- NetApp Windows Host Utility Kit 5.3
- NetApp MPIO 3.4
- SnapDrive 6.3 P2

3. Set up the server network:
   a. Network 1: Server management (ILO) (optional)
   b. Network 2: Client access (VM BRIDGE)
   c. Network 3: Live migration network (optional)
   d. Network 4: Heartbeat network
   e. Network 5: iSCSI network (as needed)
   f. Network 6: CSV network (for redirected I/O) (optional)
   g. HBA: FCP connections (as needed)
4. Set up SDW:
   a. Set up Transport Protocol defaults.
   b. Set up individual controllers.
   c. Provision the disks from the preferred controller IP address.
5. Set up the cluster:
   a. Create a Windows cluster.
   b. Enable CSVs.
   c. Use SnapDrive to set up the LUN to be used for the quorum drive.
   d. Use Cluster Manager to set up failover cluster settings for the quorum system.
6. Create clustered shared volumes:
   a. Use SnapDrive for Windows to create CSVs.
   b. Open System Manager and convert LUNs to thin-provisioned LUNs.
7. Set up SMHV:
   a. Use SnapDrive to create a single clustered drive to be used for the SnapInfo directory.
   b. Install SMHV on every node in the cluster.
   c. Add the cluster to the SMHV management console. (Refer to SMHV Installation and Administration Guide.)
   d. Create a base dataset. (Refer to SMHV Installation and Administration Guide.)

e. Create a backup policy. (Refer to SMHV Installation and Administration Guide.)

8. Create virtual machines as needed.
9. Open SMHV to add virtual machines to the appropriate datasets.
10. Repeat tasks 8 and 9 as needed.

How to Select 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 using Hyper-V.

Table 3 summarizes the pros and cons of each choice to help make the decision-making process easier.

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

Storage Container	Pros	Cons
Pass-through disk	• Delivers fastest performance • Has simplest storage path because file system on host is not involved • Has better alignment under SAN • For shared storage–based pass-through, has 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 take VM Snapshot copy • Is used exclusively and directly by a single VM • 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 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 • Has no risk of 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 thus saves space under certain circumstances • Can have compact operation to reduce	• 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

Storage Container	Pros	Cons
	physical file size	because of inherent size limitation of 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 will cause 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 or later and does not have to be installed separately.

Viewing Installed VSS Providers

To view the VSS providers installed on the host, complete 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
```

Verifying 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:

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

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

SMHV: When Virtual Machine Backups Take Too Long to Complete

If a 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 takes a hardware Snapshot copy of all the LUNs in the VM 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 KB ID: 975354 on the Microsoft Support site.

SMVH: 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, as follows:

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

   Note: All VM migrations should be performed using SCVMM.
3. Create an SMHV dataset for each CSV and make sure that all VMs that reside in that CSV are placed into that dataset. For best results, do not allow VMs owned by multiple nodes to coreside within the same CSV.
4. Create a backup policy for each dataset that matches the 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 using a Windows PowerShell script that can be executed by SMHV at the beginning of the backup job.

Performance Test Carried Out for SQL Server Virtual Machine

Configuration details are as follows:

- Guest operating system: Windows Server 2008 R2 SPI Virtual Machine
- Host operating system: 2-node W2K8R2 SP1 cluster
- Application installed on guest: SQL Server 2008
- Database layout:
  - Database aa1 on C drive (vhd0)
  - Databases bb1, bb2, and bb3 across E drive (vhd1) and F drive (vhd2) with data on E and log on F
  - vhd0, vhd1, and vhd2 are in same volume (CSV LUN)

Table 4 lists the basic test cases.

Table 4) Basic test cases.

No.	Basic Test Cases	Result
1.	Load data on aa1, bb1, bb2, bb3 (100 rows on each)	Pass
2.	Run DBCC CHECKDB on databases	Pass
3.	Back up using SMHV	Pass
4.	Load data on aa1, bb1, bb2, bb3 (added another 500 rows on each)	Pass
5.	Run DBCC CHECKDB on databases	Pass
6.	Restore vhd files (restore VM from SMHV)	Pass
7.	Run DBCC CHECKDB on databases	Pass
8.	Checked database contents (100 rows on each)	Pass

Table 5 lists the longevity test cases.

Table 5) Longevity test cases.

No.	Continuous Tests	Result
1.	Continuously load data on aa1, bb1, bb2, bb3 (up to 10,000 rows on each)	Pass
2.	While loading data, run DBCC CHECKDB on databases	Pass
3.	While loading data, back up using SMHV	Pass
4.	While loading data, run DBCC CHECKDB on databases	Pass
5.	Restore vhd files (restore VM from SMHV)	Pass
6.	Run DBCC CHECKDB on databases	Pass

SnapManager for Hyper-V 1.2 Application-Consistent and Crash-Consistent Backup Performance Numbers

Test Bed

A combination of Windows Server 2008 R2 VMs of size 10GB and Windows XP VMs of size 2GB.

- 500 VMs running in 4-node cluster
- 1,000 VMs running in 8-node cluster

Note: The durations recorded are in minutes.

Table 6 lists the SnapManager for Hyper-V 1.2 application-consistent and crash-consistent backup performance numbers.

Table 6) SnapManager for Hyper-V 1.2 application-consistent and crash-consistent backup performance numbers.

Operating System	Windows Server 2012 DC		Windows Server 2008 R2 EE	
Nodes	4-Node Cluster	8-Node Cluster	4-Node Cluster	8-Node Cluster

Operating System	Windows Server 2012 DC				Windows Server 2008 R2 EE			
Backup Type	Crash-Consistent Backup	Application-Consistent Distributed Backup	Crash-Consistent Backup	Application-Consistent Distributed Backup	Crash-Consistent Backup	Application-Consistent Backup	Crash-Consistent Backup	Application-Consistent Backup
Iteration 1	8.48	30.40	20.50	61.28	12.77	141.80	42.92	291.13
Iteration 2	8.35	29.92	21.82	63.80	12.80	139.78	43.87	319.63
Iteration 3	8.22	30.87	21.67	61.43	12.97	138.08	43.55	308.98
Iteration 4	8.25	30.51	21.89	63.48	12.98	139.50	43.89	318.60
Average	**8.33**	**30.43**	**21.47**	**62.50**	**12.88**	**139.79**	**43.56**	**309.59**

Guidelines for SMHV 1.1 on Clustered Data ONTAP 8.1.1 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:

1. Collect all VMs sharing a storage system volume in one dataset.
2. Do not start application-consistent and crash-consistent backups at the same time.
3. 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.

4. Make sure that VMs running across multiple dedicated Hyper-V hosts or clusters do not share a storage system volume.

If your system is not configured to these requirements, you might receive following error:

```
Snapshot operation not allowed due to clones backed by snapshots.
```

This message occurs when, in an environment operating in clustered Data ONTAP, SnapDrive for Windows creates a SIS clone to facilitate the Hyper-V automatic recovery process for an application-consistent backup of a virtual machine (VM). While the clone operation is in progress, no other Snapshot operation can finish on that storage system volume.

When this message occurs, SnapDrive for Windows retries the Snapshot copy creation for `_backup`.

When creating a backup Snapshot copy, SnapDrive for Windows provides the following registry entries to increase the retry attempt and interval in between each retry attempt:

- CloneMaxRetryCount (default 100 retry)
- CloneRetryInterval (default 10 seconds)

As a result, when an application-consistent backup of a dataset is in progress, any other Snapshot operation on the storage system volumes involved in the backup set might fail with the error:

```
Snapshot operation not allowed due to clones backed by snapshots.
```

References

NetApp Documents

- Data ONTAP 7.3 Block Access Management Guide for iSCSI and FC
  http://now.netapp.com/NOW/knowledge/docs/ontap/rel731_vs/pdfs/ontap/bsag.pdf
- Data ONTAP DSM 3.5 for Windows MPIO Installation and Administration Guide
  http://now.netapp.com/NOW/knowledge/docs/mpio/win/reldsm35/pdfs/install.pdf
- TR-3326: SnapMirror Sync and SnapMirror Semi-Sync Overview and Design Considerations
  http://media.netapp.com/documents/tr-3326.pdf
- TR-3437: Storage Subsystem Resiliency Guide
  http://media.netapp.com/documents/tr-3437.pdf
- TR-3446: SnapMirror Async Overview and Best Practices Guide
  http://media.netapp.com/documents/tr-3446.pdf
- TR-3450: Active-Active Controller Overview and Best Practices Guidelines
  http://media.netapp.com/documents/tr-3450.pdf
- TR-3505: NetApp Deduplication for FAS: Deployment and Implementation Guide
  http://media.netapp.com/documents/tr-3505.pdf
- TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide
  http://media.netapp.com/documents/tr-3701.pdf
- SnapManager 1.1 for Hyper-V Installation and Administration Guide
  http://now.netapp.com/NOW/knowledge/docs/smhv/relsmhv10/pdfs/install.pdf
- SnapDrive 6.4.1 for Windows Installation and Administration Guide
  http://support.netapp.com/documentation/productlibrary/index.html?productID=30049
- Windows Host Utilities 5.3 Installation and Setup Guide
  http://support.netapp.com/NOW/knowledge/docs/hba/win/relwinhu53/pdfs/setup.pdf

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 passthru 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?

14.1 NetApp Web Sites

- Configuring Alarms
  https://library.netapp.com/ecmdocs/ECMM1278625/html/opsmgr/monitor5.htm

- Getting Started with Data ONTAP PowerShell Toolkit
  https://communities.netapp.com/docs/DOC-6162
- NetApp High-Availability Solutions
  www.netapp.com/us/solutions/infrastructure/data-protection/high-availability.html
- NetApp Interoperability Matrix Tool
  http://support.netapp.com/matrix
- NetApp SnapDrive for Windows
  www.netapp.com/us/products/management-software/snapdrive-windows.html
- NetApp OnCommand Management Software
  www.netapp.com/us/products/management-software/snapdrive-windows.html
- NetApp Support (formerly NOW)
  http://support.netapp.com
- Operations Manager
  www.netapp.com/us/products/management-software/operations-manager.html
- Remote LAN Module (RLM)
  http://support.netapp.com/NOW/download/tools/rlm_fw/info.shtml

14.2 Microsoft References

- Configuring Disks and Storage
  http://technet.microsoft.com/en-us/library/ee344823(WS.10).aspx
- A Description of the Diskpart Command-Line Utility
  http://support.microsoft.com/default.aspx?scid=kb;en-us;300415
- About Virtual Machines and Guest Operating Systems
  http://technet.microsoft.com/en-us/library/cc794868(WS.10).aspx
- Configuring a SAN Environment for VMM
  http://technet.microsoft.com/en-us/library/cc764269.aspx
- Configuring Virtual Networks
  http://technet.microsoft.com/en-us/library/cc816585(WS.10).aspx
- Hyper-V: Using Hyper-V and Failover Clustering
  http://technet.microsoft.com/en-us/library/cc732181(WS.10).aspx
- Hyper-V: Using Live Migration with Cluster Shared Volumes in Windows Server 2008 R2
  http://technet.microsoft.com/en-us/library/dd446679(WS.10).aspx
- Install the Hyper-V Role on a Full Installation of Windows Server 2008
  http://technet.microsoft.com/en-us/library/cc794929(WS.10).aspx
- Install the Hyper-V Role on a Server Core Installation of Windows Server 2008
  http://technet.microsoft.com/en-us/library/cc794852(WS.10).aspx
- Microsoft Hyper-V Server 2008 Configuration Guide
  www.microsoft.com/Downloads/details.aspx?familyid=E1E111C9-FA69-4B4D-8963-1DD87804C04F&displaylang=en
- Hyper-V and VHD Performance: Dynamic vs. Fixed
  http://blogs.technet.com/b/winserverperformance/archive/2008/09/19/hyper-v-and-vhd-performance-dynamic-vs-fixed.aspx
- Microsoft KB ID: 302577
  http://support.microsoft.com/kb/302577
- Microsoft KB ID: 958184
  http://support.microsoft.com/kb/958184
- Microsoft KB ID: 974909
  http://support.microsoft.com/kb/974909

- Microsoft KB ID: 975354
  http://support.microsoft.com/kb/975354
- Microsoft KB ID: 975921
  http://support.microsoft.com/kb/975921
- Microsoft KB ID: 978157
  http://support.microsoft.com/kb/978157
- Microsoft KB ID: 979743
  http://support.microsoft.com/kb/979743
- Microsoft KB ID: 2406705
  http://support.microsoft.com/kb/2406705
- Microsoft Virtual Hard Disk FAQ
  http://technet.microsoft.com/en-us/bb738381.aspx
- New Installation of VMM
  http://technet.microsoft.com/en-us/library/cc793149.aspx
- New in Hyper-V Windows Server 2008 R2 Part 1- Dedicated Networks
  http://blogs.technet.com/b/jhoward/archive/2009/05/04/new-in-hyper-v-windows-server-2008-r2-part-1-dedicated-networks.aspx
- Performance Tuning Guidelines for Windows Server 2008 R2
  http://msdn.microsoft.com/en-us/windows/hardware/gg463392.aspx
- Planning for Disks and Storage
  http://technet.microsoft.com/en-us/library/dd183729%28WS.10%29.aspx
- Planning for Backup
  http://technet.microsoft.com/en-us/library/dd252619(WS.10).aspx
- Virtual Network Manager
  http://technet.microsoft.com/en-us/library/cc754263.aspx
- VMM System Requirements
  http://technet.microsoft.com/en-us/library/cc764328.aspx
- What's New in Windows Server 2008 R2 Hyper-V Performance and Scale?
  http://blogs.msdn.com/b/tvoellm/archive/2009/08/05/what-s-new-in-windows-server-2008-r2-hyper-v-performance-and-scale.aspx
- Windows Server 2008 R2 and Microsoft Hyper-V Server 2008 R2: Hyper-V Live Migration Overview and Architecture
  http://www.microsoft.com/en-us/download/details.aspx?displaylang=en&id=12601
- 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	January 2012	Initial release
Version 1.0.1	June 2012	Updated to include best practices for SMHV 1.1 features on clustered Data ONTAP 8.1.1
Version 1.0.2	January 2013	Updated with Windows Server 2012 support

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