



Technical Report

NetApp SnapManager 2.1 for Hyper-V on Clustered Data ONTAP 8.3 Best Practices Guide

Vinith Menon, NetApp
November 2014 | TR-4355

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

1	Executive Summary	4
2	Scope	4
3	SnapManager 2.1 for Hyper-V	4
3.1	Technical Details	4
3.2	Other Reference Documentation	5
4	SnapManager for Hyper-V Planning	5
4.1	Storage Considerations	5
4.2	Product Summary: Supported Platforms and Guidelines	5
5	SnapManager for Hyper-V Architecture	7
5.1	SMHV Port Usage	8
5.2	SMHV Architecture	8
6	MetroCluster	10
7	SnapManager for Hyper-V Backup Types	11
7.1	Application-Consistent Backup	11
7.2	Crash-Consistent Backup and Restore	13
7.3	Adding a Hyper-V Parent Host or Host Cluster	14
7.4	SMHV Backup Mechanism in Windows Server 2008 R2 SAN Environments	15
7.5	SMHV Backup Mechanism for Windows Server 2012 and 2012 R2 SAN Environments	16
7.6	SMHV 2.1 Backup Process in Windows Server 2012 and Windows Server 2012 R2 SMB 3.0 Environments	17
7.7	Scheduled Backups and Retention Policies	18
7.8	Handling Saved-State Backups of VMs	19
7.9	Quick-Migration and Live-Migration Best Practices	21
7.10	Restore Process	21
7.11	Mount a Backup	22
8	SnapManager for Hyper-V Disaster Recovery	24
8.1	Get-VMsFromBackup Cmdlet	24
8.2	Prerequisites	24
8.3	Fail Over VMs to Secondary Site	24
8.4	Fail Back VMs to the Primary Site	26
9	SMHV Disaster Recovery for VMs in Hyper-V over SMB Environments	27
9.1	Disaster Recovery Workflow on Secondary Site for SMB Shares	27
9.2	Set File-Level or Directory-Level ACLs or DACLs for Access-Denied Issues Related to SMB Share VM Restore	28

10 SnapVault Integration.....	29
References.....	30
NetApp Knowledge Base Articles	30

LIST OF TABLES

Table 1) SMHV product summary.	6
Table 2) SMHV terminology.	7

LIST OF FIGURES

Figure 1) SMHV architecture.	8
Figure 2) SMHV deployed to manage virtual entities in a clustered Data ONTAP environment.	9
Figure 3) SnapInfo settings.	10
Figure 4) ScopeSnapshots PowerShell command.	12
Figure 5) ScopeSnapshots setting in Registry Editor tool.....	13
Figure 6) Backup Dataset wizard.	14
Figure 7) Hyper-V infrastructure and associated storage during an application-consistent SMHV backup.	15
Figure 8) SMHV backup process for Windows Server 2012 and Windows Server 2012 R2 SAN environments.	17
Figure 9) SnapManager 2.1 for Hyper-V architecture.....	18
Figure 10) SnapVault integration in SMHV.....	29
Figure 11) SMHV SnapVault options and Snapshot labels.	30

1 Executive Summary

Server virtualization is a major component of data center virtualization, and it 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 NetApp clustered Data ONTAP® 8.3 architecture. It describes best practices for using NetApp SnapManager® for Hyper-V (SMHV), a NetApp tool that uses NetApp Snapshot® technology for backup, recovery, and replication of virtual machines (VMs) in a Hyper-V environment.

NetApp has been on the forefront of solving complex business problems with its innovative technology breakthroughs and end-to-end solutions approach. This technical report is not intended to be a definitive implementation or solutions guide. Expertise might be required to solve specific deployments. Contact your local NetApp sales representative to speak with one of our Microsoft Hyper-V solutions experts. NetApp is dedicated to providing data availability to keep your applications running using the least amount of resources while leveraging your data protection investment to accelerate your organization.

2 Scope

This document provides prescriptive guidance and best practices for deploying SnapManager for Hyper-V in a Windows® virtual environment. Unless otherwise specified, the best practices described in this document apply to SnapManager 2.1 for Hyper-V installed with NetApp SnapDrive® 7.1 for Windows (SDW) in a clustered Data ONTAP 8.3 environment.

3 SnapManager 2.1 for Hyper-V

With the adoption of virtualization technologies, data centers have been transformed, and the number of physical servers has been 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 arrangement can create performance issues, especially while the virtual environment is being backed up. That is because many VMs use host network and CPU resources concurrently. Therefore, backups that once completed during nonbusiness hours might no longer be able to finish within their backup window of time.

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

Backups, restores, and disaster recovery (DR) can place a demanding 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 their associated datasets. The following sections detail the best practices for deploying and using SnapManager for Hyper-V.

3.1 Technical Details

SMHV offers the following capabilities:

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

- Integrates with Microsoft Hyper-V Volume Shadow Copy Service (VSS) writer to quiesce the Hyper-V VMs before creating an application-consistent Snapshot copy of the VM in SAN LUNs
- Supports backup and restore of VMs running on continuously available SMB shares that are hosted on systems based on Data ONTAP 8.3

Note: Backup operations are performed by using a remote VSS plug-in located in Data ONTAP.

- Allows administrators to create application-consistent backups of Hyper-V VMs if Microsoft Exchange, Microsoft SQL Server®, or any other VSS-aware application is running on virtual hard disks (VHDs) in the VM
- Provides replication (through NetApp SnapMirror® technology) and vaulting (through NetApp SnapVault® technology) of backup sets to secondary locations for DR planning
- Supports backup and restore of shared VMs configured through Windows Failover Clustering (WFC) for high availability (HA) and also on Microsoft Cluster Shared Volumes (CSVs)

Note: SMHV supports seamless processing of scheduled VM backups, regardless of any VM failovers.

- Supports management of multiple remote Hyper-V parent systems from a single console
- Supports fast, crash-consistent backup and restore of VMs

3.2 Other Reference Documentation

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

- [TR-4172: Microsoft Hyper-V over SMB 3.0 with Clustered Data ONTAP: Best Practices](#)
- [SnapManager 2.1 for Hyper-V Installation and Administration Guide](#)
- [SnapDrive 7.1 for Windows Installation Guide](#)

4 SnapManager for Hyper-V Planning

This section summarizes storage considerations, supported platforms, and guidelines for deploying SMHV.

4.1 Storage Considerations

SMHV supports backup and restore of VMs on dedicated disks, CSVs, or SMB 3.0 shares. It backs up only VM data stored in VHDs that reside on NetApp storage. It does not back up data on pass-through or direct-attached iSCSI or virtual Fibre Channel (vFC) disks.

SMHV does not support master boot record LUNs for VMs running on shared volumes or CSVs. It does support LUNs created on thin-provisioned volumes, and it can perform backups and restores on these volumes.

To host VMs in SMB 3.0 shares in Windows Server 2012 and Windows Server 2012 R2, the storage system should run clustered Data ONTAP 8.2 or later versions.

SnapDrive 7.1 for Windows must be installed on the host system. NetApp recommends using SnapDrive to provision LUNs or shares to host VMs.

4.2 Product Summary: Supported Platforms and Guidelines

Table 1 lists the new features for SMHV 2.1 and the supported versions of SDW and Data ONTAP.

Table 1) SMHV product summary.

SMHV Version	New Features	Supported SnapDrive Version	Supported Data ONTAP Versions
SMHV 2.1	Hyper-V over SMB in Windows Server 2012 and Windows Server 2012 R2	SDW 7.1	<ul style="list-style-type: none"> • Data ONTAP 8.3 (clustered Data ONTAP) • Data ONTAP 8.2 (7-Mode and clustered Data ONTAP) • Data ONTAP 8.1.X (7-Mode and clustered Data ONTAP)

Note: To see the full support matrix, refer to the [NetApp Interoperability Matrix Tool \(IMT\)](#).

SMHV Deployment Guidelines

Consider the following guidelines before deploying SMHV:

- For this version of SMHV, you must install the supported SnapDrive version indicated in Table 1.
- After you upgrade to SMHV 2.1, it is not possible to revert to the previous version of SMHV.
- To use SMHV in Hyper-V over SMB environments, you must have VMs hosted in systems running clustered Data ONTAP 8.2 or later versions.
- Before upgrading to the latest SMHV, you must upgrade to the corresponding SnapDrive version.
- Before upgrading to the latest SnapDrive version, you must upgrade the dependent components, such as the Data ONTAP Device-Specific Module (DSM) and the Windows Host Utilities Kit. For information about supported versions of the DSM and the Windows Host Utilities Kit, refer to the [SnapDrive 7.1 for Windows documentation](#) or to the [IMT](#).
- SDW is required on Hyper-V parent hosts but not on client hosts. For Windows Server Failover Clustering configurations, SDW and SMHV must be installed on each node of the cluster.
- For backup and restore of VMs in SMB 3.0 shares (Hyper-V over SMB 3.0), SnapDrive 7.1 for Windows and SnapManager 2.1 for Hyper-V are required.
- Remote installation is supported for standalone and cluster nodes in a domain. SMHV cannot be installed on a host that is part of another domain.

Best Practice

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

License Requirements

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

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

The following licenses are required for systems running clustered Data ONTAP 8.2 or later versions:

- SnapRestore license
- FCP or iSCSI license (for SAN environments)
- CIFS license (for Hyper-V over SMB)
- NetApp FlexClone[®] license (for Hyper-V over SMB)

SnapMirror and SnapVault licenses are optional.

Platform Support

For current information, refer to the [IMT](#).

5 SnapManager for Hyper-V Architecture

Table 2 defines key terms used throughout this document.

Table 2) SMHV terminology.

Term	Description
Dataset	A dataset is a group of VMs that helps to protect data by using retention, scheduling, and replication policies. Datasets can be used to group VMs that have the same protection requirements. A VM can be a member of multiple datasets, which can be useful for VMs that belong to multiple groups (for example, a VM running the SQL Server instance for a Microsoft Office SharePoint® Server [MOSS] configuration might have to belong to both the SQL Server and the MOSS datasets). A dataset cannot contain a mix of VMs hosted on both SMB shares and SAN LUNs.
Protection policies	Policies make it possible to schedule or automate the backups of datasets at a predefined time. For example, retention policies provide the capability to schedule retention periods for older backups, replication policies make it possible to replicate block changes to the SnapMirror destination volume after the VM backup is created, and SnapVault policies enable the vaulting of Snapshot copies to a SnapVault destination. Policies include other capabilities that allow scripts to be run 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 by using SnapMirror relationships or by updating the Snapshot copies to a SnapVault destination. Backups are performed on the whole dataset, which is a logical collection of VMs, with the option of updating the SnapMirror relationship and the SnapVault destination as part of the backup on a per-job basis. Similarly, restores can be performed at an individual VM level.
Backup retention policy	Retention policies can be used to specify how long to keep a dataset backup, based on either a specified time or the number of backups. Policies can be created to specify the retention period, allowing administrators the flexibility to meet varying service-level agreements (SLAs) in their environment.
Alert notification	Alert notifications are created on a per-scheduled-backup-job basis and are sent by e-mail to administrator-defined accounts. An alert notification can be configured to e-mail the specified account after every backup, although NetApp does not recommend doing this because the number of e-mails can become unmanageable. Configuring alerts to notify administrators after an error or warning within a backup offers a more useful and practical alert level.
Unprotected resources	Unprotected resources are VMs that are not part of any dataset. These resources can be protected by adding them to a dataset.
Application-consistent backup and restore	Application-consistent backups are created in coordination with VSS so that the applications running in the VM are quiesced before a Snapshot copy is created. Such a backup guarantees the integrity of application data and therefore can be safely used to restore the VM and the applications running in the VM to a consistent state.
Crash-consistent backup	In crash-consistent backups, the state of data is equivalent to what would be found following a catastrophic failure that abruptly shuts down the system. The data in the backup is the same as it would be after a system failure or a power outage. This type of backup is much faster than other types. A restore from such a backup is equivalent to a reboot following an abrupt shutdown.

5.1 SMHV Port Usage

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

- **808**, the SMHV and SDW default port
- **4094**, if SDW is configured to use the HTTP protocol
- **4095**, if SDW is configured to use the HTTPS protocol

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

5.2 SMHV Architecture

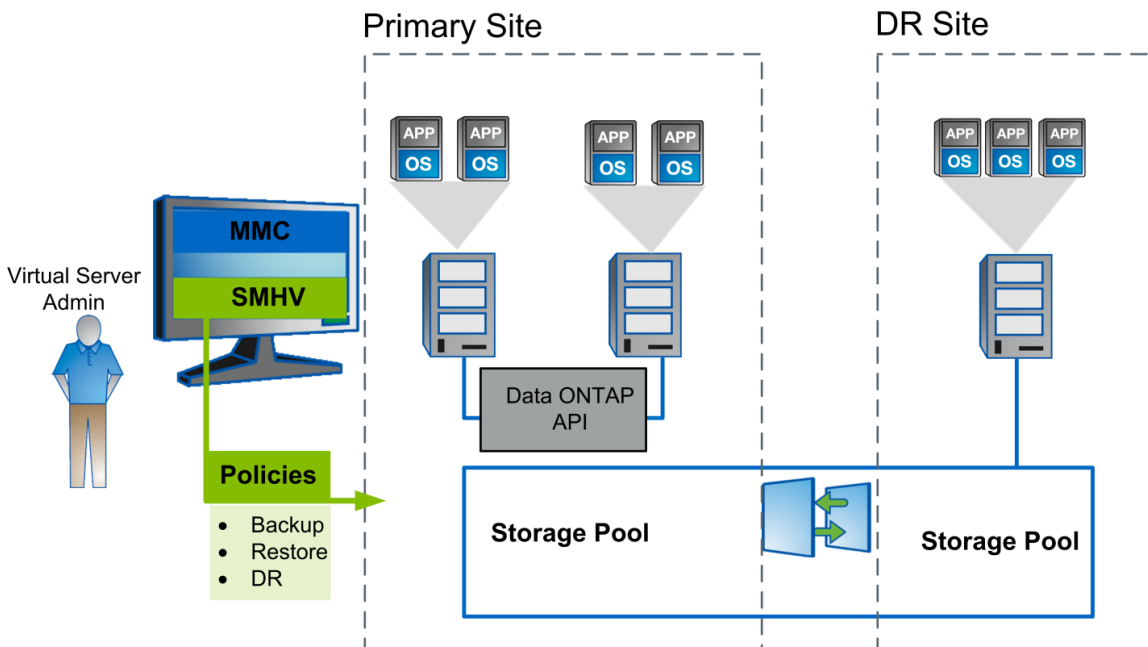
SMHV must be installed on the Hyper-V parent host to create the backup and restore of the VMs running in the Hyper-V parent host. It provides a management console based on Microsoft Management Console (MMC) and a Windows PowerShell® snap-in for management operations. SMHV allows multiple Hyper-V parent hosts to be managed from a single console, which can be installed on a Hyper-V parent host, on other Windows (non-Hyper-V) servers, or on client systems such as Windows 8.

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

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

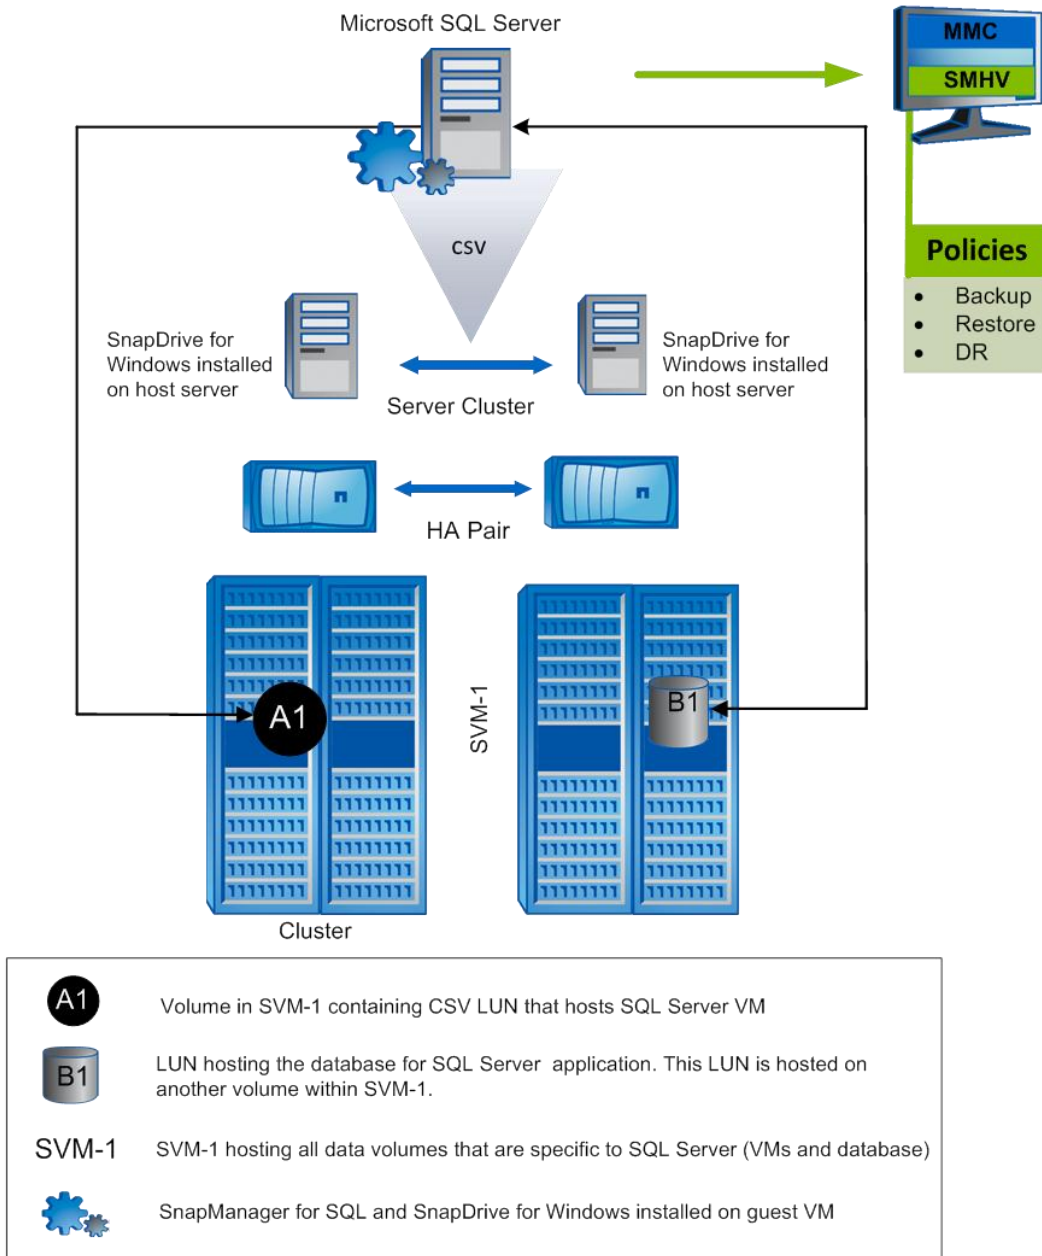
Figure 1 illustrates the SMHV architecture and the components that work together to provide a comprehensive and powerful backup and recovery solution for Hyper-V environments.

Figure 1) SMHV architecture.



With clustered Data ONTAP systems, customers can host all VMs that pertain to a workload within volumes mapped to a NetApp storage virtual machine (SVM, formerly called Vserver). This ability simplifies provisioning and protection management. Figure 2 shows how SMHV can be deployed to manage virtual entities in a SAN environment with clustered Data ONTAP 8.2 or later versions.

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



SMHV Components

SnapManager for Hyper-V includes three major components:

- SMHV SnapInfo directory
- SMHV report settings
- SMHV event notifications

SMHV SnapInfo Directory

The SMHV SnapInfo directory stores the backed-up VM's metadata. This folder can be set up by specifying the SnapInfo settings in the Hosts Management wizard, as shown in Figure 3. 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 in that host can be added to a dataset. SMHV also creates a

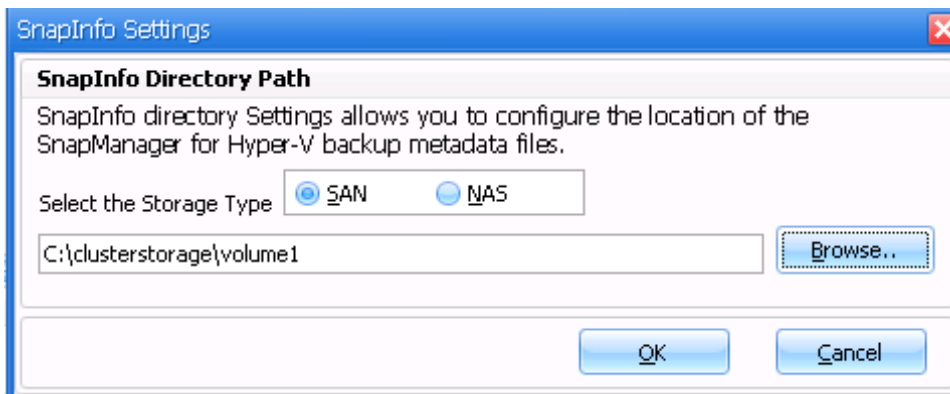
Snapshot copy of the SnapInfo directory after the backup is completed. The naming convention for the SnapInfo Snapshot copy is `smhv_snapinfo_hostname_timestamp`.

With SnapManager 2.1 for Hyper-V, SnapInfo can be hosted on CSV LUNs, SMB 3.0 shares, or dedicated LUNs.

Note: SnapInfo can be hosted on SMB 3.0 and CSV LUNs only with systems using Windows Server 2012 and Windows Server 2012 R2.

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

Figure 3) SnapInfo settings.



Best Practice

For clustered Data ONTAP 8.3, NetApp recommends locating the SnapInfo directory on a share in a separate volume in the SVM and not making it part of the other data volumes. For example, if SMHV is protecting SQL Server VMs that are hosted in a CSV LUN in a volume, then the user must make sure that the SnapInfo directory is not part of this volume of the SVM. This separation simplifies VM restoration and disaster recovery.

SMHV Report Settings

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

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

SMHV Event Notifications

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

6 MetroCluster

NetApp MetroCluster™ provides both local failover (HA) and site failover capabilities with clustered Data ONTAP by providing RAID-level replication of data between sites. For Microsoft Hyper-V environments in which a Hyper-V replica is leveraged, using MetroCluster results in the creation of multiple copies of the VM data. NetApp recommends considering making duplicate or multiple copies of the VM data. MetroCluster is ideally suited for clustered Hyper-V deployments.

For more information, refer to the [MetroCluster](#) documentation.

7 SnapManager for Hyper-V Backup Types

This section describes the types of backup allowed by SMHV.

7.1 Application-Consistent Backup

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

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

VSS coordinates Snapshot copy–based backup and restore. It includes these additional components:

- **VSS requestor.** The VSS requestor is a backup application, similar to the SMHV application or to 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 these tasks:

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

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

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

SMHV integrates 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. Using VSS hardware provider for Data ONTAP, it coordinates the backup operation to create a consistent Snapshot copy. SMHV enables the creation of 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. Applications in the VM are restored to the state that existed at the time of the backup. SMHV restores the VM to its original location.

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

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

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

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

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

For detailed information about VM backup, refer to [Planning for Backup](#) on the [Microsoft TechNet site](#).

Note: Application-consistent backup of SAN VMs through SMHV 2.1 might fail with the error message `Either the specified VM(s) are not present or they cannot be backed up online.` (Under Backup Options, do not select the Allow Saved-State VM Backup checkbox.)

After checking the basic information and all of the event viewer logs from Hyper-V and the application event logs in the VM that is triggered for backup, SMHV registers the following event in the event log: Event ID: 13 Windows cannot perform an online backup of the system because scoped snapshots are enabled.

To resolve this problem, disable the use of scoped snapshots by creating the following registry value on the server.

```
PATH: HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\SystemRestore\  
DWORD: ScopeSnapshots  
Value: 0
```

You can indicate this setting by executing a basic PowerShell command:

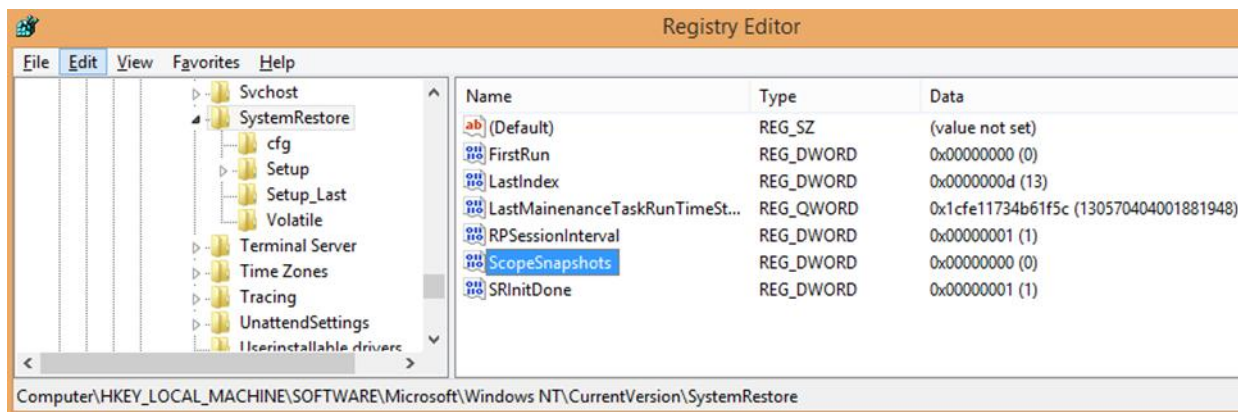
```
New-Itemproperty -path 'HKLM:\Software\Microsoft\Windows NT\CurrentVersion\SystemRestore' -Name  
ScopeSnapshots -value 0
```

Figure 4 shows the `ScopeSnapshots` PowerShell command in the command-line interface, and Figure 5 shows the `ScopeSnapshots` setting in the Windows Registry Editor tool.

Figure 4) ScopeSnapshots PowerShell command.

```
PS C:\> New-Itemproperty -path 'HKLM:\Software\Microsoft\Windows NT\CurrentVersion\SystemRestore' -Name ScopeSnapshots -value 0  
  
ScopeSnapshots : 0  
PSPath          : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion\SystemRestore  
PSParentPath    : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\Software\Microsoft\Windows NT\CurrentVersion  
PSChildName     : SystemRestore  
PSDrive         : HKLM  
PSProvider      : Microsoft.PowerShell.Core\Registry
```

Figure 5) ScopeSnapshots setting in Registry Editor tool.



Note: To complete the backup successfully, set the registry entry as recommended by Microsoft and restart the backup job.

Best Practices

- For better performance with SMHV backups, split the VMs into multiple datasets.
- When you create a dataset, NetApp recommends that you select all of the VMs that reside on a particular Data ONTAP LUN. This approach enables you to fit all of the backups into one Snapshot copy and reduce the space consumption on the storage system.

7.2 Crash-Consistent Backup and Restore

Backups created through SMHV can be either application consistent or crash consistent. Application-consistent backups are created in coordination with Microsoft VSS so that the applications running in the VM are quiesced before the Snapshot copy is created. Such a backup assures the integrity of application data; therefore, it can be safely used to restore the VM and the applications running in the VM to a consistent state.

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

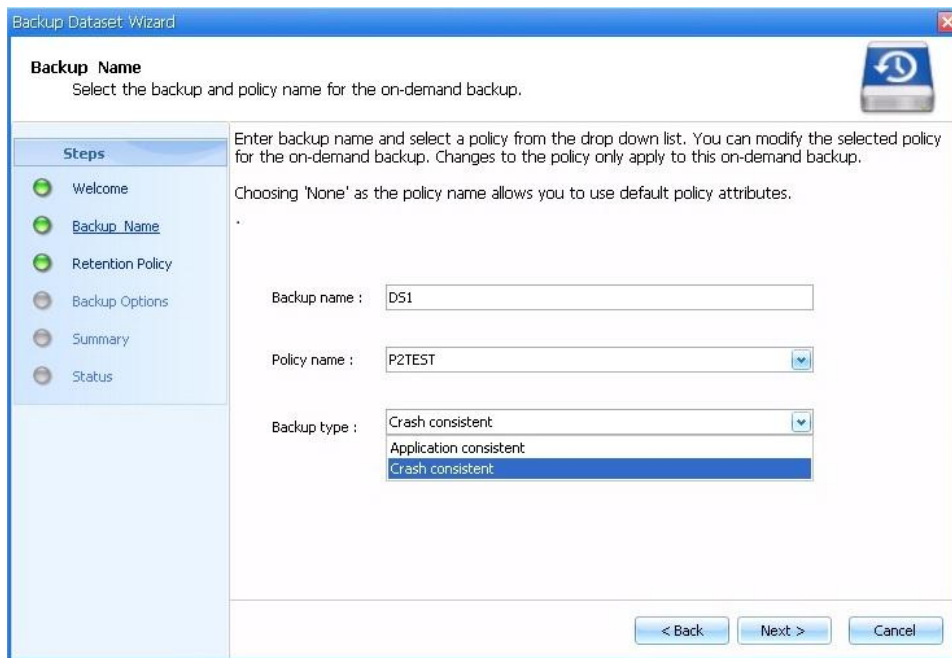
- Application-consistent backups are slower because of VSS involvement with the parent and guest OS. Because both the application writer in the VM and the Hyper-V writer in the parent OS are involved, failure to back up any of the components causes the backup process to fail.
- Hyper-V writer uses the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

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

A crash-consistent backup of a VM does not use VSS to quiesce data, and it does not result in autorecovery. This backup simply creates a Snapshot copy on the NetApp storage system for all of the LUNs used by the VMs involved in the dataset. The data in the backup is the same as it would be after a system failure or a power outage. All SMHV functions, such as scheduling, restore, script execution, SnapMirror updates, backup retention, and so on, are supported for crash-consistent backups as well. Crash-consistent backups are also supported both for SAN and SMB 3.0 environments.

Figure 6 shows the Backup Dataset wizard with both the application-consistent and the crash-consistent backup types displayed.

Figure 6) Backup Dataset wizard.



- Note:** The saved-state backup policy is not applicable to crash-consistent backup and restore because crash-consistent backups do not involve the Hyper-V VSS writer.
- Note:** The SMHV backup policy name can contain only alphanumeric characters. The SMHV GUI checks for this condition and blocks the policy creation if nonalphanumeric characters are included. This check is not present in the `Create-SMHVPolicy` cmdlet, so creating the SMHV policy with this cmdlet allows the user to create the policy, but subsequent backups using this policy will fail.
- Note:** SMHV supports parallel-execution, crash-consistent, and application-consistent backups. It also supports parallel crash-consistent backup execution. Because of a timeout error in the underlying SnapDrive for Windows software, however, users might observe some issues while such operations are executed.
- Note:** Restore of crash-consistent backups in SMB environments fails if the directories hosting them are renamed after the backup is performed.

Best Practices

- The crash-consistent backup feature is not a replacement for application-consistent backups. It enables the user to have frequent recovery points and therefore to have frequent crash-consistent backups and fewer application-consistent backups.
- The crash-consistent backup feature can be used to create the latest backup of all of the data just before performing an application-consistent restore operation of a VM.

7.3 Adding a Hyper-V Parent Host or Host Cluster

If a single host is added, SMHV manages the dedicated VMs on that host. If a host cluster is added, SMHV manages the shared VMs on the host cluster. If there is a plan to add a host cluster, SMHV must be installed on each cluster node. SMHV 2.1 supports remote installation of SMHV on all nodes of the Windows cluster from a single node.

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

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

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

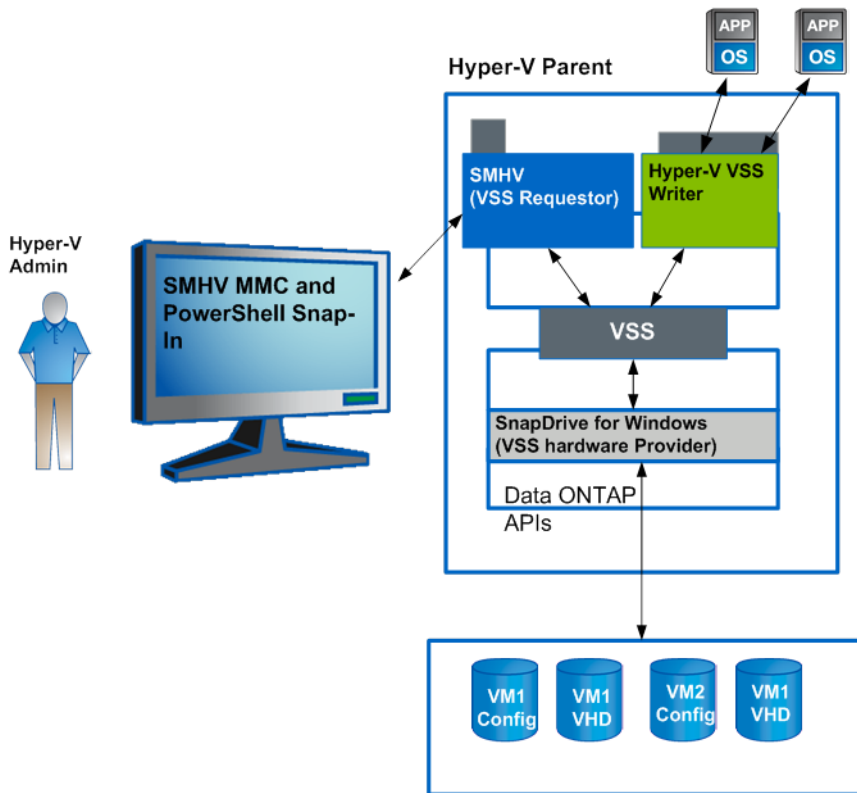
7.4 SMHV Backup Mechanism in Windows Server 2008 R2 SAN Environments

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

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

Figure 7 shows a high-level overview of the typical SMHV architecture on the primary site storage, where the backup process takes place in a SAN environment.

Figure 7) Hyper-V infrastructure and associated storage during an application-consistent SMHV backup.



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 concurrent backups occur, 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. For more information, refer to the [Microsoft Support](#) website.

Best Practice

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

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

Note: The grouping of VMs hosted on SMB shares and SAN LUNs in a single dataset is not supported.

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

Best Practices

- When creating a dataset, select all VMs that reside on a particular Data ONTAP LUN. This makes it possible to have all backups in one Snapshot copy and reduce 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, verify that the backup schedules of these datasets do not overlap.
- If a VM Snapshot copy location is changed to a different Data ONTAP LUN after the VM is created, create at least one VM Snapshot copy by using Hyper-V Manager before creating a backup by using SMHV. If this is not done, the backup might fail.
- For clustered Data ONTAP systems, all VMs related to an application can be isolated to a single NetApp SVM. A single dataset can be created to back up and restore these VMs, which simplifies the backup, restore, and DR processes.

7.5 SMHV Backup Mechanism for Windows Server 2012 and 2012 R2 SAN Environments

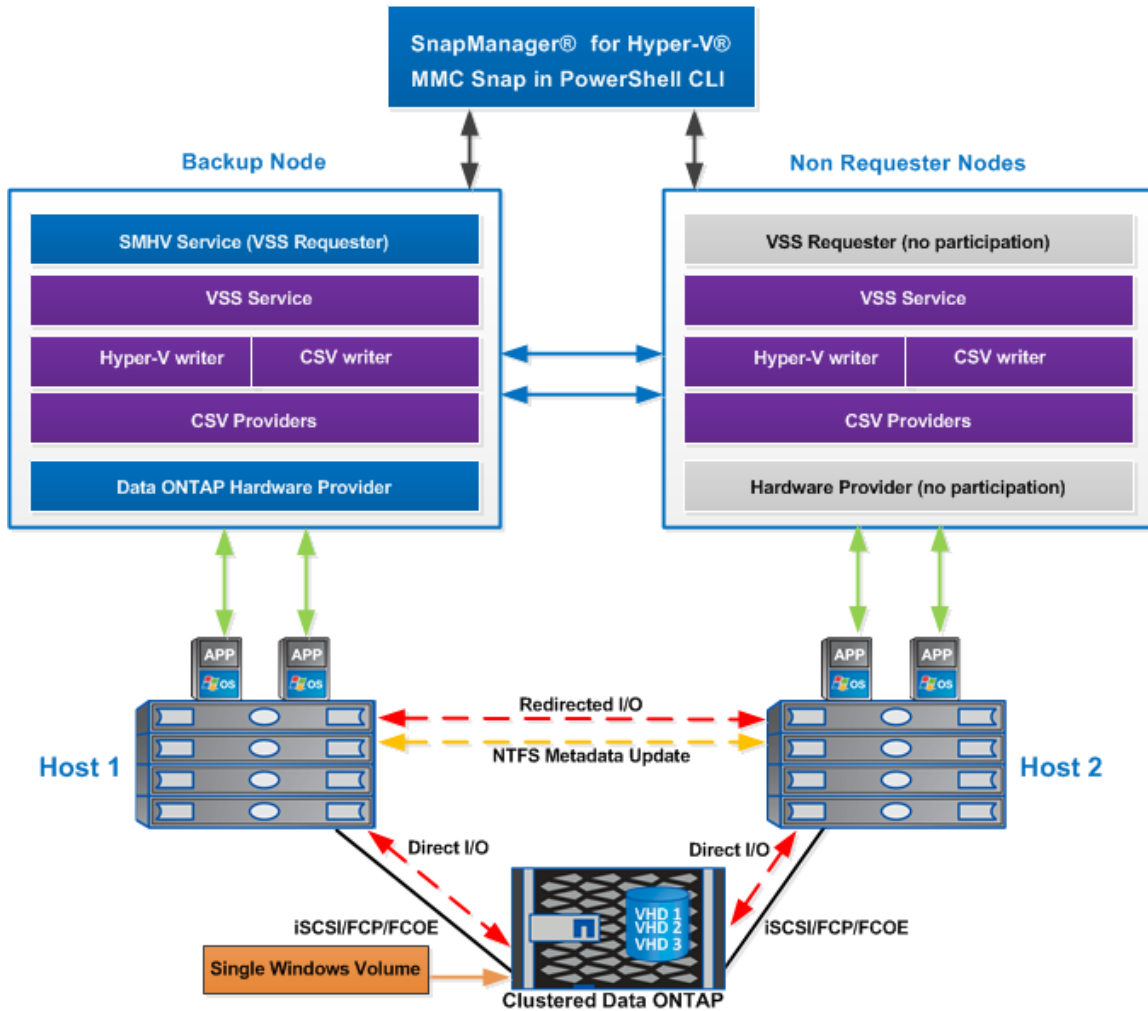
In Windows Server 2012 and 2012 R2, Microsoft introduced the CSV proxy file system (CSVFS), which provides a cluster shared storage LUN with a single and consistent file namespace while still using the underlying NTFS file system. In Windows Server 2012 and Windows Server 2012 R2, the CSVs appear as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

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

In Windows Server 2012 and Windows Server 2012 R2, CSVFS introduced distributed application-consistent backups. This feature allows backup of all of the VMs in a cluster to be consistent in one single application-consistent backup. To achieve this distributed backup mechanism, Microsoft has introduced a new CSV writer and CSV provider.

Figure 8 illustrates the SMHV backup process for Windows Server 2012 and Windows Server 2012 R2.

Figure 8) SMHV backup process for Windows Server 2012 and Windows Server 2012 R2 SAN environments.



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

Distributed application-consistent backup is also more space efficient because it creates only one Snapshot copy for each volume instead of creating one Snapshot copy for each node and volume combination. The space saving is significant if large numbers of nodes are involved in the backup. In addition, Data ONTAP limits the maximum number of Snapshot copies that can be stored for a volume, so this enhancement allows more backups to be stored for a VM.

Note: SMHV does not support hosting the VMs in such CSVs on asymmetric clusters in Windows Server 2012 and Windows Server 2012 R2. SMHV supports BitLocker functionality for CSVs provisioned through SDW. VMs can be hosted in encrypted CSVs in Windows Server 2012 and Windows Server 2012 R2.

Note: SMHV does not support backup of VMs that have LUNs attached through vFC in the VM. However, during the VM backup, the LUN presented to the VM is not backed up.

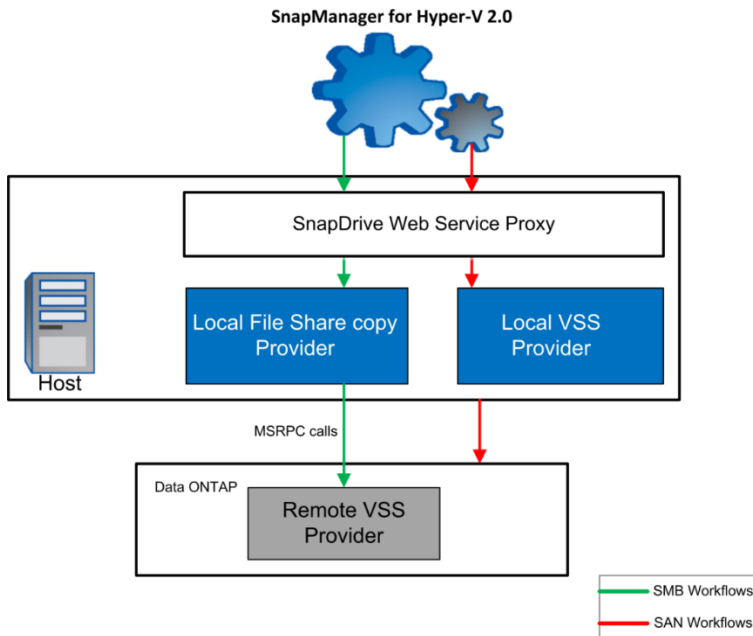
7.6 SMHV 2.1 Backup Process in Windows Server 2012 and Windows Server 2012 R2 SMB 3.0 Environments

Clustered Data ONTAP 8.3 supports two important features developed specifically for Windows Server 2012 environments: continuously available shares for Hyper-V over SMB and remote VSS. Users can

create continuously available SMB shares by using the provisioning templates in SnapDrive 7.1 for Windows and host VMs on them. These VMs can be backed up with SMHV by using remote VSS. Users can initiate a backup from either the SMHV GUI or PowerShell cmdlets.

When the VSS requestor (SMHV) adds an SMB 3.0 share containing Hyper-V VMs to the VSS Snapshot set, VSS invokes the new SMB File Share Copy Provider component in Windows Server 2012 and Windows Server 2012 R2 to send the Microsoft Remote Procedure Call (MSRPC) commands to the SMB target (the storage system) to coordinate the VSS backups. The new File Share Shadow Copy Agent (the remote VSS provider) running on the SMB target is responsible for creating the hardware Snapshot copy. Clustered Data ONTAP 8.3 implements the File Share Shadow Copy Agent (remote VSS hardware provider) to perform the application-consistent backup copy of the SMB shares. Figure 9 shows these workflows.

Figure 9) SnapManager 2.1 for Hyper-V architecture.



Note: Hyper-V Server 2012 and Hyper-V Server 2012 R2 do not support the use of application-consistent backup for VMs in SMB shares.

7.7 Scheduled Backups and Retention Policies

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

Backup Scheduling

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

Best Practices

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

Retention Policies

The following list describes the retention tags available in SMHV:

- **Hourly.** Hourly intervals
- **Daily.** A specified time within a 24-hour period
- **Weekly.** A specified day and time within a 7-day period
- **Monthly.** A specified day and time within a calendar month
- **Unlimited.** Never-deleted backups

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

NetApp recommends using the retention policies to meet specific SLAs and also to maintain a supported number of NetApp Snapshot copies on the underlying volumes. For SMHV, one backup creates two Snapshot copies on the storage systems for data consistency. (For more information, refer to [KB ID: 2010607](#).) For example, setting a retention policy of 30 backups on an hourly backup limits to 60 the maximum number of Snapshot copies associated with the backup. However, if the retention policy is configured as 30 days, the Snapshot copy limit per volume is reached in 5 days, and backups fail from that point on.

Best Practice

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

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

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

7.8 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 is shut down, an online backup cannot be performed. In some cases, VMs are in the saved state or are shut down for maintenance, but backups must still proceed, even if an online backup is not possible. To make such backups possible, the VMs that are in the saved state or are shut down can be moved to a different dataset with a policy that allows saved-state backups.

Note: Selecting the Allow Saved-State VM Backup checkbox enables SMHV to back up the VM by using the saved state. If this option is selected, SMHV does not fail the backup when the Hyper-V

VSS writer backs up the VM by using the saved state or performs an offline backup of the VM. Performing a saved-state or offline backup might cause downtime. For more information about online and offline VM backups, refer to [Planning for Backup](#) on the [Microsoft TechNet site](#).

Note: The Allow Saved-State Policy option is not applicable for crash-consistent backups because the VM is backed up regardless of the state.

PowerShell Cmdlets

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

- `Add-SMHVDataSet`
- `Add-SMHVHost`
- `Add-SMHVPolicy`
- `Delete-Backup`
- `Get-Backup`
- `Get-SMHVDataSet`
- `Get-SMHVHost`
- `Get-SMHVPolicy`
- `Get-VMsFromBackup`
- `Invoke-SMHVConfigureHost`
- `Invoke-SMHVRemoteHostInstall`
- `Invoke-SMHVRemoteHostUninstall`
- `New-Backup`
- `Remove-SMHVDataSet`
- `Remove-SMHVHost`
- `Remove-SMHVPolicy`
- `Restore-Backup`
- `Set-SMHVDataSet`
- `Set-SMHVPolicy`

Use the `Get-Help` command for each cmdlet to learn about their functioning.

7.9 Quick-Migration and Live-Migration Best Practices

Best Practices

- SMHV cannot back up a VM that is actively undergoing migration. If 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.
- To improve not only the success rate of the backups but also the overall VM performance, NetApp recommends migrating VMs only when a significant gain in performance can be achieved.
- Avoid SMHV-related operations during storage live migration because such operations could corrupt the VM.
- After VM storage is migrated from one LUN or share to another, the VM can no longer be restored with the previous Snapshot copy. As a safety net, create an SMHV backup immediately after VM storage migration is complete.
- After VM storage is migrated from one share to another hosted on a different volume, the SnapMirror and SnapVault relationship must be verified. Also, previously existing Snapshot copies cannot be restored from the SnapVault storage system.

7.10 Restore Process

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

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

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

Aside from these differences in restore types, the SMHV uses the following process flow during a restore:

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

During the restore, the following warnings 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.
- The VM to be restored is currently running on the host, and 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.

If the VM no longer exists, it can still be restored if the LUNs on which it was created still exist. The LUNs must have the same drive letters and Windows volume globally unique identifiers (GUIDs) that they had at the time of backup.

If the VM no longer exists, it can still be restored by the selection of a backup to which it belonged. If it was removed from all datasets before it was deleted, the user can still restore it by selecting Unprotected Resources and then selecting from the list a backup to which it belonged.

Best Practice

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

Note: SMHV does not back up the cluster configuration of the VM, so it does not restore the cluster configuration. If the VM and the cluster configuration are lost, the VM can be restored from SMHV, but it must be manually made highly available. For more information, refer to [Failover Cluster Step-by-Step Guide: Configuring a Two-Node File Server Failover Cluster](#).

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

Note: Windows Server 2008 does not support restoring deleted VM crash-consistent backups, but Windows Server 2012 does support it.

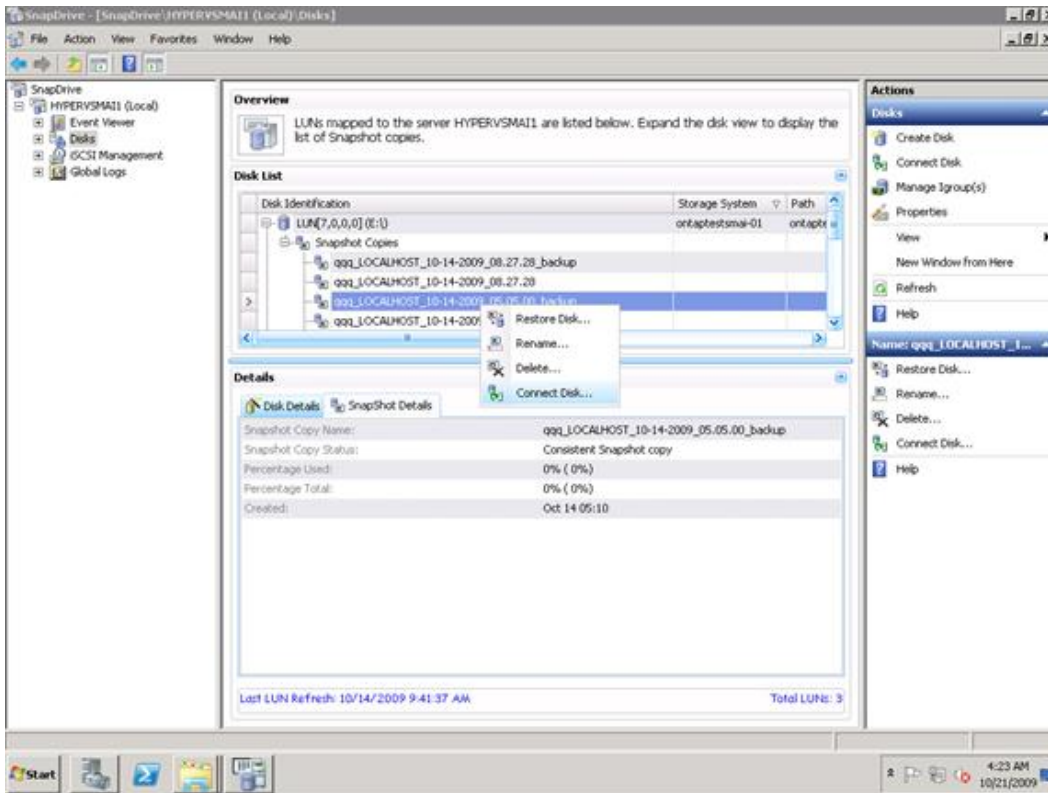
Note: If a crash-consistent backup was created by earlier versions of SMHV, it cannot be restored with the latest version of SMHV.

7.11 Mount a Backup

SnapDrive for Windows can be used to mount backups. The mounted backup is a clone of the protected VM. After it is mounted, the backup is displayed in the explorer of the Hyper-V host and can be browsed.

To mount a backup by using SDW, complete the following steps:

1. Select the LUN and, from the list of 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. In the Connect Disk wizard, click Next.
3. If the LUN is a dedicated disk, proceed to step 4. If the LUN is a Windows cluster resource, perform the following step in the Specify Microsoft Cluster Services Group pane. Select only one of the following actions and then click Next:
 - To use an existing cluster group, select it from the Group Name drop-down list.
 - To create a new cluster group, select the Create a New Cluster Group option.

Note: When selecting a cluster group for the LUNs, choose the cluster group that the application will use. If a volume mount point is being created, the cluster group is already selected because the cluster group owns the root volume physical disk cluster resources. NetApp recommends creating new shared LUNs outside of the cluster group.
4. To use CSVs, select the Add option.
5. In the Select LUN Properties pane, either select a drive from the list of available drive letters or enter a mount point for the LUN that is being connected. 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 pane, choose an initiator for the LUN.
7. In the Select Initiator Group Management pane, specify either automatic or manual igroup management.
8. On the Completing the Connect Disk page, verify all of the settings and click Finish.

Note: To change any settings, click Back to return to the previous wizard pages.
9. Browse the backup by selecting the drive letter on the explorer of the Hyper-V host.

8 SnapManager for Hyper-V Disaster Recovery

Users can fail over and fail back Hyper-V VMs by using Windows PowerShell cmdlets in the SMHV PowerShell option. To use this feature, the Windows PowerShell cmdlet `restore-backup` must be used along with the `-RestoreToAlternateHost` switch and the server name. For example:

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

8.1 Get-VMsFromBackup Cmdlet

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

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

8.2 Prerequisites

The following prerequisites must be completed before failover can take place:

- There must be a site A (the primary site) containing storage systems and a standalone Hyper-V host system or a Hyper-V host cluster. VMs running on these hosts reside on NetApp storage.
- There must be a site B (the secondary site) containing storage systems and a Hyper-V host or cluster (the same as that of the primary site).
- SDW and SMHV must be installed on both site A and site B.
- A SnapMirror relationship must be initialized from site A to site B.
- A Hyper-V host or cluster on site A must be added to SMHV, and the VMs must be backed up by using SMHV. The policy to update SnapMirror after backup must be selected so that, after each backup, the secondary site is updated with new Snapshot copies of VMs and the SnapInfo directory.

8.3 Fail Over VMs to Secondary Site

To fail over VMs to the secondary site for both SAN and SMB DR scenarios, complete the following steps:

1. Connect to all of the LUNs from the secondary storage system volumes.

Note: If the secondary system is a cluster, go to the node where the cluster group is online and connect to all of the LUNs from that node in the cluster. The LUN type and mount point must be the same as those of the primary system. SDW breaks the SnapMirror relationship and also performs a SnapRestore operation. If the volume contains only one LUN, SDW performs a volume-based SnapRestore operation, and the SnapMirror relationship is then in an uninitialized state. If the volume contains multiple LUNs, SDW performs a single-storage-system SnapRestore operation, and the SnapMirror relationship is broken off.

2. Restore the SnapInfo LUN from the last Snapshot copy created by SMHV.
3. Add the secondary host or cluster in SMHV and configure it with the SnapInfo path.
4. Use the `Get-VMsFromBackup` cmdlet to display the list of VMs present in backup metadata.


```

PS C:\> Get-VMsFromBackup

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

```

5. Use the `Get-Backup` cmdlet to display the backups for each VM.

```

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

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

```

Note: For other VMs that have multiple VHDs or Snapshot copies, you might not be able to view all of the files in the `FilesList`.

```

PS C:\Users\administrator.SDDEV> Get-Backup -Resourcename DR-VHD2

BackupName      : ds-2VHD_01-15-2013_16.25.16
FilesList       : {\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Virtual
                  Machines\582E6C41-DC55-4798-A48A-91B4930C1224.xml,
                  \\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Virtual
                  Machines\582E6C41-DC55-4798-A48A-91B4930C1224\*,
                  \\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Virtual Hard
                  Disks\DR-VHD2_AC9C9DC0-954F-44DD-807E-EE030D67B450.avhdx,
                  \\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Virtual Hard
                  Disks\DR-VHD2_4DD5894D-C455-4389-A75F-4C0D0FA8A1B1.avhdx...}
RetentionType   : hourly
DatasetName     : ds-2VHD
BackupId        : ds-2VHD_01-15-2013_16.25.16
BackupTime      : 1/15/2013 4:25:16 PM
BackupType      : Application consistent
BackedupVMs    : {DR-VHD2}

```

6. For these types of VMs, you can display the complete `FilesList` by using a variable and then accessing the `FilesList` object.

```

PS C:\Users\administrator.SDDEV> $v = Get-Backup -Resourcename DR-VHD2
PS C:\Users\administrator.SDDEV> $v[0].FilesList
\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Virtual Machines\582E6C41-DC55-4798-A48A-91B4930C1224.xml
\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Virtual Machines\582E6C41-DC55-4798-A48A-91B4930C1224\*
\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Virtual Hard Disks\DR-VHD2_AC9C9DC0-954F-44DD-807E-EE030D67B450.avhdx
\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Virtual Hard Disks\DR-VHD2_4DD5894D-C455-4389-A75F-4C0D0FA8A1B1.avhdx
\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Virtual Hard Disks\DR-VHD2_E614C558-2FB1-4A89-9AE0-313388F219B5.avhdx
\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Virtual Hard Disks\DR-VHD2.vhdx
\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\VHD-NUM2\New Virtual Hard Disk_BF007965-B4BE-4F2A-AA5F-C097202F1588.avhdx
\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\VHD-NUM2\New Virtual Hard Disk_D9582C5F-35C9-4A44-BE41-A60F480B8B5C.avhdx
\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\VHD-NUM2\New Virtual Hard Disk_61D63E48-971B-49FB-8650-DDCAE146FA0F.avhdx
\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\VHD-NUM2\New Virtual Hard Disk.vhdx
\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Snapshots\38C69B0A-C5D0-42FD-A3B8-A903084B0ECA.xml
\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Snapshots\38C69B0A-C5D0-42FD-A3B8-A903084B0ECA.*
\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Snapshots\65B79880-844F-432D-A4E4-0C7D1BD569C6.xml
\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Snapshots\65B79880-844F-432D-A4E4-0C7D1BD569C6.*
\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Snapshots\7D2C697F-8E32-46FA-BE60-7FE073459538.xml
\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Snapshots\7D2C697F-8E32-46FA-BE60-7FE073459538.*

```

7. Use the `Restore-Backup` cmdlet with the VM GUID that was obtained from step 4 by using the `Get-VMsFromBackup` cmdlet. Use the `-RestoreToAlternateHost` switch and specify the secondary host or cluster name as `-server` parameter.

```

PS C:\Users\administrator.S0DEV> Restore-Backup -server clab-a13-15 -ResourceId 224DDECE-5C50-401F-8B48-9797EBD58CC6 -BackupName ds-new_01-14-2013_14_20_48 -RestoreToAlternateHost -Verbose -DisableVerifySnapshot -VirtualMachinePath "\\172.17.175.81\vol5_share\vm2013" -SnapshotFilePath "\\172.17.175.81\vol5_share\vm2013" -VHDs @(@{"SourceFilePath" = "\\172.17.175.81\vol2_share\vm2013\Virtual Hard Disks\vm2013.vhdx"; "DestinationFilePath" = "\\172.17.175.81\vol5_share\vm2013\Virtual Hard Disks\vm2013.vhdx"})
VERBOSE: Starting restore-backup
VERBOSE: Processing restore-backup..
VERBOSE: The input parameters are being validated.
VERBOSE: Validating the Input Parameters ....
VERBOSE: Validating the Input VM
VERBOSE: Validating the Input backup
VERBOSE: DR - Found All parameters for DR
VERBOSE: DR - Virtual Machine Path validated
VERBOSE: DR - Snapshot Paths validated
VERBOSE: DR - VHD Paths validated
VERBOSE: DR - Validated all NAS Restore parameters
VERBOSE: The input parameters validation successful.
VERBOSE: Performing operation "restore-backup" on Target "clab-a13-15".
VERBOSE: Proceeding with restore-backup .The Administrator confirmed the Input parameters.
VERBOSE: Invoking Restore Backup
VERBOSE: Source Type : Virtual Machine Configuration Path ; Dest Type : Virtual Machine Configuration Path
Source : \\172.17.175.81\vol2_share\vm2013\Virtual Machines\224DDECE-5C50-401F-8B48-9797EBD58CC6.xml ;
Dest : \\172.17.175.81\vol5_share\vm2013\Virtual Machines\224DDECE-5C50-401F-8B48-9797EBD58CC6.xml
-----
VERBOSE: Source Type : Virtual Machine Configuration Directory Path ; Dest Type : Virtual Machine Configuration Directory Path
Source : \\172.17.175.81\vol2_share\vm2013\Virtual Machines\224DDECE-5C50-401F-8B48-9797EBD58CC6\* ;
Dest : \\172.17.175.81\vol5_share\vm2013\Virtual Machines\224DDECE-5C50-401F-8B48-9797EBD58CC6\*
-----
VERBOSE: Source Type : Virtual Machine Hard Disks ; Dest Type : Virtual Machine Hard Disks
Source : \\172.17.175.81\vol2_share\vm2013\Virtual Hard Disks\vm2013.vhdx ;
Dest : \\172.17.175.81\vol5_share\vm2013\Virtual Hard Disks\vm2013.vhdx
-----

```

Note: If the secondary system is a cluster, make sure that the LUNs on which VMs reside are online on the cluster node that owns the cluster group. For a VM that contains one or more Snapshot copies, the `-SnapshotFilePath` parameter is mandatory.

Note: For a VM that contains more than two VHDs, you must provide an array of VHDs in the following format:

```

-VHDs @(@{"SourceFilePath" = "\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\DR-VHD2\Virtual Hard Disks\DR-VHD2.vhdx"; "DestinationFilePath" = "\\172.17.175.81\vol5_share\DR\Dr-2VHD\DR-VHD2\Virtual Hard Disks\DR-VHD2.vhdx" }, @{"SourceFilePath" = "\\172.17.175.81\srcvol2_share\DR\Dr-2VHD\VHD-NUM2\New Virtual Hard Disk.vhdx"; "DestinationFilePath" = "\\172.17.175.81\vol5_share\DR\Dr-2VHD\VHD-NUM2\New Virtual Hard Disk.vhdx"})

```

Note: Each SourceFilePath VHD in the backed-up VM must have a corresponding DestinationFilePath specified as an array of hash tables, as shown in the previous example.

If the secondary site is an active site with its own VM LUNs and SnapInfo LUN, then to restore the VMs present in the primary site to the secondary site, complete the following steps:

1. Connect the primary SnapInfo LUN to the secondary host by breaking the mirrored volume.
2. Perform a SnapRestore operation from the last SMHV SnapInfo Snapshot copy.
3. Copy the contents to the already-existing SnapInfo copy on the secondary site.

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

8.4 Fail Back VMs to the Primary Site

To fail back VMs to the primary site, complete the following steps:

1. Transfer the data from the secondary storage system back onto the primary storage system, selecting the appropriate method for your situation:
 - If the primary site was completely destroyed, new storage must be provisioned. Then the user must initialize the SnapMirror relationship from the secondary site to the primary one (because this is a new relationship) to get the data back. After the relationship is initialized and the data is back on the primary site, this relationship can be released.
 - If the primary site was down temporarily, the user must transfer to it only those changes that happened on the secondary site while the primary one was gone. To do this, resync the existing SnapMirror relationship in the reverse direction (resync from the secondary to the primary site).

2. After the data on the secondary site is synchronized with the data on the primary site, go to the SnapDrive user interface on the secondary site and initiate a SnapMirror update for each of the LUNs on the secondary site. If this is not done, SDW uses the SMHV backup Snapshot copy to restore the LUNs on the primary site during the resyncing in step 1. The LUN in the backup Snapshot copy is actually a LUN clone, so this outcome must be avoided by forcing one more SnapMirror update.

Note: Creating an SMHV backup (with the Update SnapMirror checkbox selected) from the secondary site has the same effect as manually performing the SnapMirror update from the SDW GUI. Most users prefer to use the SMHV backup instead of manually performing a mirror update because the SMHV backup can be scripted, whereas updating the mirror is a tedious job and prone to user error (such as forgetting to update a LUN).
3. Connect to all LUNs on the primary site (with the same type and the same mount points). If the primary site is a cluster, go to the node where the cluster group is online and connect to all of the LUNs from that node in the cluster. If a resync in reverse direction has been performed, a new broken (or uninitialized) SnapMirror relationship exists from the secondary to the primary site. This relationship should be released.
4. Restore the SnapInfo directory from its last Snapshot copy created by SMHV.
5. Add the primary host or cluster in SMHV MMC and configure it with the SnapInfo path.
6. Use the `Get-VMsFromBackup` cmdlet to display a list of VMs present in backup metadata.
7. Use the `Get-Backup` cmdlet to display the backups for each VM.
8. Use the `Get-VMsFromBackup` cmdlet with VM GUID (from step 6) and the `Get-Backup` cmdlet (from step 7). Use the `-RestoreToAlternateHost` switch and specify the primary host or cluster name as `-server` parameter. If the primary site 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.
9. If the primary site is the cluster, make the VMs highly available by using the failover cluster UI or the Windows PowerShell cmdlet.

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

1. Select the method of powering off and refreshing the secondary site that is appropriate for your situation:
 - If the secondary site is a standalone host, shut it down and delete the VMs running on it. Disconnect the SnapInfo directory and the disks containing VMs that use SnapDrive.
 - If the secondary site is a cluster, offline the VM resource and the VM configuration resource for all of the VMs. Delete these resources from the cluster, delete all VMs from Hyper-V Manager, and disconnect all disks that use SnapDrive.
2. Resync the SnapMirror relationship from the primary to the secondary site.

9 SMHV Disaster Recovery for VMs in Hyper-V over SMB Environments

When the secondary site's SMB shares must be brought up, additional steps are required. This section describes the tasks involved in disaster recovery in environments that use Hyper-V over SMB protocol.

9.1 Disaster Recovery Workflow on Secondary Site for SMB Shares

In clustered Data ONTAP, the SnapMirror destination volume is created with its type designated as data protection (DP). Clustered Data ONTAP does not allow creating junction paths when volumes are created. Therefore, SMB 3.0 shares cannot be created on the SnapMirror or Snap vault destination for DR restores.

In the case of SMHV restore for SMB backups, the share on the secondary volume must be exposed to allow the Hyper-V servers to access the VM configurations and data on the secondary site. After SnapMirror updates, quiesces, or breaks the SnapMirror relationship, the following tasks must be completed for a share to be created on the secondary volume:

1. The R/W volume must be unmounted.
2. The R/W volume must be mounted on a junction path.
3. A new SMB share must be created on the destination R/W volume and exposed to the clients.

9.2 Set File-Level or Directory-Level ACLs or DACLs for Access-Denied Issues Related to SMB Share VM Restore

To set the access-control lists (ACLs) at the VM directory level, a new file-directory policy should be created or the `Filer Security` command should be run.

After the VM is failed over to the secondary site, the VM directory in the primary site might point to the primary domain account or computer account. When the VM is restored in the secondary site, the user cannot browse to the share location and might receive the `access-denied` message. This problem can be eliminated by setting the SMB ACLs or discretionary ACLs (DACLs) on a per-file or per-directory basis with granular access control by using the file-directory setting `Fsecurity`.

`Fsecurity` is a new NetApp WAFL[®] (Write Anywhere File Layout) security feature for setting Data ONTAP file-directory security through a policy-based infrastructure that uses the Data ONTAP command-line interface and ZAPIs. This feature makes it possible to create an end-to-end DR solution in an SMB environment by controlling the ACLs at the VM-directory level when restoring the VM from the secondary site.

To set the file or directory levels for ACLs or DACLs, complete the following steps:

1. Create a new NTFS security identifier (SID) targeting the SVM (formerly Vserver).

```
f2552-187-44-45::> ntfs create -vserver SMHV_SVM -ntfs-sd ntfspermissions -owner
virtualcloud\administrator
(vserver security file-directory ntfs create)
```

Note: The SID in this example is named `ntfspermissions`.

2. Use the `ntfs show` command to view the existing security descriptors.

```
f2552-187-44-45::> ntfs show
(vserver security file-directory ntfs show)

Vserver: SMHV_SVM

NTFS Security      Owner Name
Descriptor Name    -----
ntfspermissions    VIRTUALCLOUD\Administrator
```

3. Set the SMB ACLs or DACLs on a per-file or per-directory basis by using granular access control.

```
f2552-187-44-45::> dacl add -vserver SMHV_SVM -ntfs-sd ntfspermissions -access-type allow -
account virtualcloud\administrator -rights full-control
(vserver security file-directory ntfs dacl add)
```

4. Create a new policy to set the ACLs for specified directories or files.

```
f2552-187-44-45::> vserver security file-directory policy task add -vserver SMHV_SVM -policy-
name pntfspermissions -path /vol_SMHV_CIFS -security-type ntfs -ntfs-mode propagate -ntfs-sd
ntfspermissions
```

5. View the existing policy and the expanded denial-of-service properties.

```
f2552-187-44-45::> vserver security file-directory policy task show -vserver SMHV_SVM -path
/vol_SMHV_CIFS
```

```
Vserver: SMHV_SVM
Policy: pntfspermissions

Index   File/Folder      Security      NTFS      NTFS Security
-----  -
1       /vol_SMHV_CIFS  ntfs         propagate ntfspermissions
```

6. Apply the policy that was created.

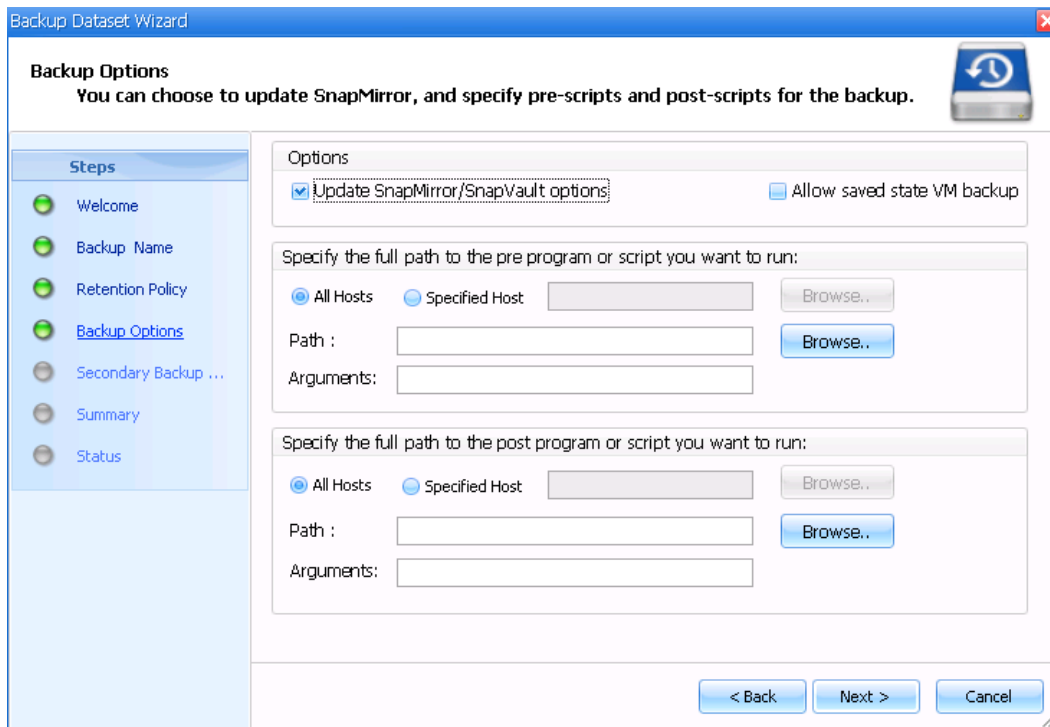
```
f2552-187-44-45::> vserver security file-directory apply -vserver SMHV_SVM -policy-name
pntfspermissions
```

After this procedure is executed, SMHV can restore the VM from the required application-consistent or crash-consistent backup.

10 SnapVault Integration

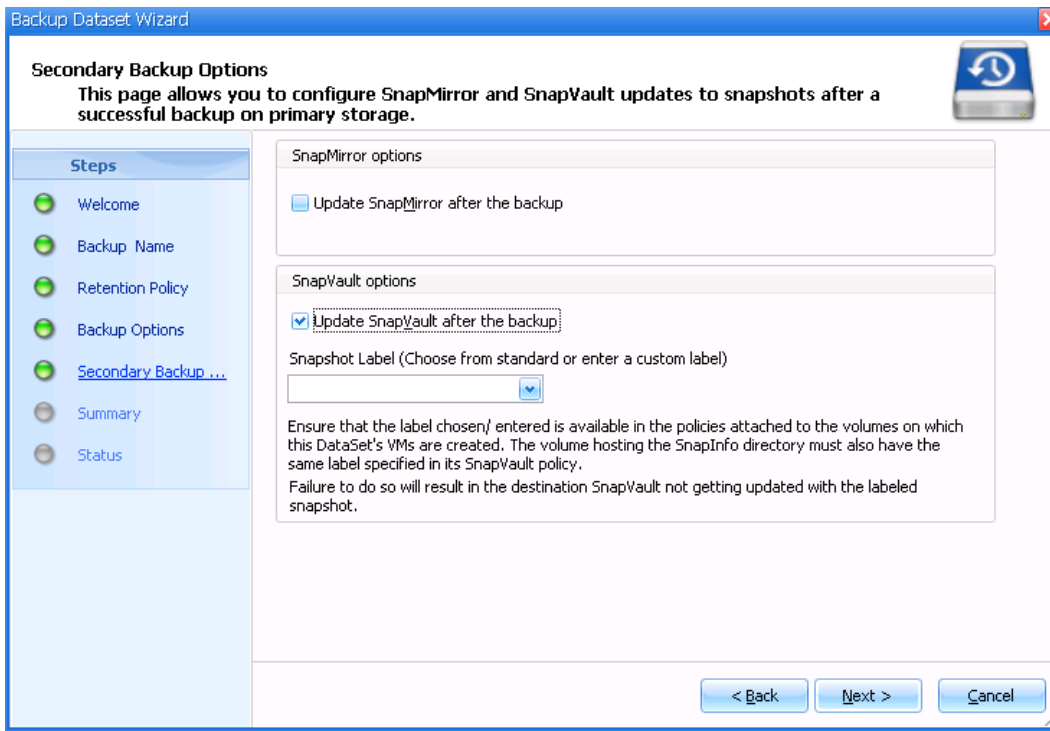
As Figure 10 shows, SMHV 2.1 supports SnapVault in clustered Data ONTAP 8.3 environments. After SMHV completes a backup operation for a dataset, the backup can be updated to the SnapVault destination.

Figure 10) SnapVault integration in SMHV.



You can also configure the retention period of a Snapshot copy in the SnapVault storage system by selecting the Labels option while configuring the SMHV backup workflow, as shown in Figure 11.

Figure 11) SMHV SnapVault options and Snapshot labels.



You can select from a set of preconfigured labels, such as `smhv_hourly`, `smhv_daily`, `smhv_weekly`, and `smhv_monthly`, or you can create a custom label. The label must be selected on the storage system and attached to the policy of the SnapVault relationship between the primary and the secondary volumes.

After the dataset backup is created, the primary backup is labeled with the Snapshot label and updated to the SnapVault destination by using SnapMirror updates. Each Snapshot update to the SnapVault destination is tied to a version universally unique identifier (UUID). This version UUID, along with other information such as the SnapVault status and label details, is stored in the SnapInfo metadata.

Note: SnapVault restore by using SMHV is not supported.

Best Practice

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

VSS requires that the provider commit a Snapshot copy within 10 seconds. If this time limit is exceeded, the Data ONTAP VSS Hardware Provider logs Event ID 4364. This limit might be exceeded because of a transient problem. If this event is logged for a failed backup, retry the backup.

References

NetApp Knowledge Base Articles

- [KB ID: 1010146](#): SMHV: How to manually restore a Hyper-V virtual machine from a Snapshot backup
- [KB ID: 1011587](#): How to migrate a Hyper-V VM to support SnapManager for Hyper-V backup
- [KB ID: 1010887](#): SMHV: How to set up SnapInfo logical unit number (LUN)
- [KB ID: 2010607](#): SMHV: Creation of two Snapshot copies for every backup
- [KB ID: 2010899](#): SMHV: Backups fail for Hyper-V virtual machines containing pass-through or iSCSI in guest disks

- [KB ID: 2014900](#): SnapManager for Hyper-V backup sets that contain Windows XP fail
- [KB ID: 2014905](#): SnapManager for Hyper-V backups fail to complete even though all virtual machines are located on NetApp LUNs
- [KB ID: 2014928](#): SMHV: During backup of CSV, hosts report NO_DIRECT_IO_DUE_TO_FAILURE
- [KB ID: 2014933](#): SMHV: Cluster Shared Volume goes offline after backup
- [KB ID: 3011206](#): SMHV: Can SnapManager 1.0 for Hyper-V exclude virtual hard disks from backups?
- [KB ID: 302577](#): How to use the Sysprep tool to automate successful deployment of Windows XP
- [KB ID: 958184](#): Virtual machine backup operations fail in Windows Server 2008 when Hyper-V virtual machine files are saved on a volume that is mounted on a failover cluster by using a volume GUID
- [KB ID: 974909](#): The network connection of a running Hyper-V virtual machine is lost under heavy outgoing network traffic on a Windows Server 2008 R2–based computer
- [KB ID: 975354](#): A Hyper-V update rollup is available for Windows Server 2008 R2
- [KB ID: 975921](#): You may be unable to perform certain disk-related operations after an exception when a hardware provider tries to create a snapshot in Windows Server 2008 R2 or in Windows 7
- [KB ID: 978157](#): MPIO removes pseudo-LUN paths for external storage devices on Windows Server 2008 SP2–based servers or on Windows Server 2008 R2–based servers
- [KB ID: 979743](#): You cannot use an MPIO storage device after a failover operation in Windows Server 2008 or in Windows Server 2008 R2
- [KB ID: 2406705](#): Some I/O requests to a storage device fail on a fault-tolerant system that is running Windows Server 2008 or Windows Server 2008 R2 when you perform a surprise removal of one path to the storage device
- [KB2770917](#): Windows 8 and Windows Server 2012 update rollup: November 2012

Note: This is a Windows Server 2012 KB fix for the following error: `Error: VSS Requestor - Backup Components failed. Writer Microsoft Hyper-V VSS Writer involved in backup or restore encountered a retryable error. Writer returned failure code 0x800423f3. Writer state is 8. This issue is caused by inclusion of direct-attached iSCSI LUNs or pass-through disks in the VSS backups.`

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.

Copyright Information

Copyright © 1994–2014 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NetApp, the NetApp logo, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States and/or other countries. A current list of NetApp trademarks is available on the Web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-4355-1114

