NetApp
Go further, faster

Technical Report

SnapManager 1.0 for Hyper-V Best Practices

Amrita Das, Ravi B, NetApp
December 2009 | TR-3805

LEVERAGING NETAPP DATA ONTAP FOR HYPER-V BACKUP, RESTORE, AND DISASTER RECOVERY

Backups, restores, and disaster recovery can place a huge overhead on the Hyper-V™ virtual infrastructure. NetApp® SnapManager® for Hyper-V 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 virtual machines (VMs) and the associated data sets. This document details the best practices for deploying and using SnapManager 1.0 for Hyper-V.

1 INTRODUCTION

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

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

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

1.1 PURPOSE AND SCOPE

The purpose of this technical report is to provide best practices for deploying SMHV to back up and recover Hyper-V VMs. It describes 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 for Hyper-V Installation and Administration Guide.

1.2 INTENDED AUDIENCE

This document is intended for Hyper-V administrators, storage administrators, backup administrators, and architects implementing a backup, restore, and disaster recovery solution for Hyper-V environments running on NetApp storage. Readers should ideally have a solid understanding of the architecture, administration, and backup and recovery concepts within a Hyper-V environment and should consider reviewing the following:

- Data ONTAP 7.2 or 7.3 System Administration Guide
- SnapManager 1.0 for Hyper-V Installation and Administration Guide
- NetApp and Hyper-V Storage Best Practices
- SnapDrive 6.2 for Windows Installation and Administration Guide

2 SNAPMANAGER 1.0 FOR HYPER-V

2.1 TECHNICAL DETAILS

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

3 PLANNING

Microsoft Windows 2008 Server R2 with Hyper-V role enabled offers various storage infrastructure configurations and provisioning methods. Refer to TR-3702, NetApp and Microsoft Virtualization Storage Best Practices, to determine the most appropriate choices for your environment.

3.1 STORAGE CONFIGURATION

Infrastructure virtualized with Microsoft Hyper-V offers support for the use of either direct-attached storage (DAS) connected through FC / iSCSI or shared storage connected to the Windows 2008 R2 Server. However, the High-Availability and Disaster Recovery features of Hyper-V integrated with Windows Failover Cluster (WFC) are only available on Shared Storage.

3.1.1 SHARED STORAGE

The traditional shared storage is a LUN provisioned from a NetApp storage system connected to multiple Hyper-V server nodes, which are a part of a single WFC. However, out of these nodes, only one Hyper-V server can perform read and write operations on the shared storage disk. This implies that when a VM is migrated from one Hyper-V host to another in a WFC (using Live Migration or Quick Migration), read/write access to the LUN with VM is transferred to the Hyper-V host to which the VM is migrated. A consequence of this is that VM migration moves all VMs that reside on that LUN to the destination Hyper-V host, since that is the only host with read/write access to that LUN. To achieve granular migration of VMs, each LUN can only contain a single VM.

Best Practice
The recommendation while deploying Hyper-V on traditional shared LUNs is to have one VM per LUN configured. All the VHDs relative to a single VM (VM with multiple drives) can reside on a single LUN provisioned as a shared storage to a WFC.

3.1.2 CLUSTER SHARED VOLUMES

A CSV is a physical disk that is a shared LUN which is accessible to multiple Hyper-V hosts, which are a part of a single WFC that can be written to and read from by these host nodes simultaneously. This is a clustering feature available to users from Windows Server 2008 R2 release, which is supported only for the Hyper-V role. CSVs can be created using the Cluster Manager User interface and from the shared disks provisioned from a NetApp storage system.

CSVs were primarily introduced to serve as a clustered file system so that the available nodes within a cluster could perform read-write operations concurrently. Hence CSV configuration within a Hyper-V environment can be used to set up highly available VMs using the live migration feature. However, configuring CSV is not mandatory to set up live migration since the traditional shared storage can also be used to achieve the objective. In such a case, it is required to adhere to the Microsoft recommendation of maintaining a one-to-one mapping of VMs to shared storage.

3.2 STORAGE LAYOUT

It is required to plan and consider the storage requirements for Virtual Machines in Microsoft Hyper-V environments ahead of the storage provisioning activities. The data drives of the virtual machines hosted on a Hyper-V server can reside on Virtual Hard Disks (VHD files), pass-through disks, or directly attached LUNs using the Microsoft iSCSI software initiator in the VMs. Regardless of the storage option, the VMs' disks are formatted with the file system native to the guest OS.

3.2.1 VIRTUAL HARD DISK FILES

VHDs are a commonly used storage option for VM configurations. With VHDs, the actual data of the VM is kept within a VHD file stored on a LUN formatted with NTFS which is connected to the Hyper-V host system. The LUNs can be accessed over Fibre Channel or iSCSI. Hyper-V supports three types of VHDs: Fixed-Size, Dynamically Expanding and the Differencing VHDs.

FIXED-SIZE VHD

In this type of VHD, the entire amount of storage space configured for the Virtual Machine is pre-allocated and hence the file doesn't expand while the VM is under operation. Out of the available VHD types, fixed-

size has the least performance overhead. For production environments, this type of VHD is most recommended.

DYNAMICALLY EXPANDING VHD

As the name indicates, this type of VHD expands while the VM is in operation as the data gets populated. The VHD expands to the maximum configuration size defined during the VHD creation. Since dynamic expansion of the disk occurs while the VM is under active I/O, this type of VHD has the highest performance overhead. This type of VHD residing on a NetApp LUN is mostly suited for a test and development environment wherein the disk performance is not critical.

DIFFERENCING VHD

Differencing VHDs are created as part of the Hyper-V snapshot process (which is not the same as NetApp Snapshot technology). Differencing VHDs function very similar to the Dynamically Expanding VHD and point to a parent VHD which can be of any type (Fixed or Dynamically Expanding). The performance impact of this type of VHD is similar to that of Dynamically Expanding and hence should follow the same recommendations when used

3.2.2 PASS-THROUGH DISKS

Pass-Through disks are physical disks presented to the Hyper-V host which are directly attached to the guest OS as raw disks. These raw disks are formatted with the guest OS file system. This disk type is used in case of large data sets and intense I/O requirements.

3.2.3 DIRECTLY ATTACHED STORAGE TO THE GUEST OS

In this method, the LUN is directly attached using the initiator software from the guest OS and is formatted with the guest OS file system. These LUNs can only be data drives as Microsoft doesn't support booting a guest OS over an iSCSI LUN for Hyper-V.

Information
SMHV can only backup VM data stored in VHDs that reside on NetApp storage. It doesn't backup data on pass-through or a direct attached iSCSI Disks. SMHV does not support MBR LUNs for VMs running on shared volumes or CSVs.

3.3 THIN PROVISIONING

The typical storage provisioning practices followed by most of the storage administrators is to pre-allocate the disk space requested by the end user. NetApp software solutions facilitate the feature of storage virtualization where the administrators can over commit on the space allocated to the users. This method, referred to as thin provisioning, allows users to utilize storage on demand. Here the available storage resources are treated as a single shared resource pool, and the consumption is accounted as it gets utilized. SMHV supports LUNs created on flexible volumes and performs backups/restores on these volumes.

3.4 DEDUPLICATION

Data deduplication is one of the flagship features of NetApp storage systems which helps eliminate duplicate data at the block level in the environment deployed on Hyper-V. This feature can be introduced in the storage systems without affecting the administration practices or tasks from the Hyper-V server end. Deduplication runs on the storage system at scheduled intervals without affecting the resources of the server hardware. Infrastructure that uses fixed-size VHDs and multiple VMs created out of a single VHD golden copy for a VDI environment can take advantage of the deduplication feature to achieve effective storage utilization.

For further information, refer to the Deduplication Implementation and Best Practices Guide.

4 SMHV SIMPLIFIES BACKUP AND RECOVERY

4.1 PREREQUISITES

SnapManager 1.0 for Hyper-V needs SnapDrive® 6.2 for Windows (SDW 6.2) to be installed as a prerequisite. 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).

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

4.2 TERMINOLOGY

DATASETS

A dataset is a grouping of virtual machines that helps you to protect data using retention, scheduling, and replication policies. You can use datasets to group VMs that have the same protection requirements. A VM could be a member of multiple datasets. This can be useful for VMs that belong to multiple groupings (e.g. a VM running the SQL Server instance for a Microsoft Office SharePoint Server configuration may need to belong to both the SQL Server and MOSS datasets).

PROTECTION POLICIES

Policies allow customers to schedule/automate the backups of the datasets at a predefined time (schedule policy), allow customers to provide retention capabilities for older backups (retention policy), and allow customers to replicate the block changes to the SnapMirror® destination volume after the VM backup is created (replication policy). Policy includes other capabilities that allow customers to run scripts before and after the backup.

BACKUP AND RECOVERY

SMHV provides local backup and recovery capability with the option to replicate backups to a remote storage system using SnapMirror relationships.

Backups are performed on the whole dataset, which is a logical collection of VMs, with the option of updating the SnapMirror relationship as part of the backup on a per job basis. Similarly, restores can be performed at an individual VM level.

BACKUP RETENTION POLICY

Retention policies can be used to specify how long you want to keep a dataset backup based on either time or number of backups Policies can be created specifying the retention period, allowing administrators flexibility to meet varying service-level agreement (SLA) levels within their environment.

ALERT NOTIFICATION

Alert notifications are created on a per scheduled backup job basis and are sent by e-mail to administrator-defined accounts. Alert notification can be configured to e-mail the specified account after every backup, although this is not recommended as the number of e-mails can become unmanageable. Configuring alerts to notify administrators after an error or warning within a backup offers a more useful and practical alert level.

UNPROTECTED RESOURCES

Unprotected resources are VMs that are not part of any dataset. These resources can be protected by adding them to a dataset.

4.3 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 HTTP protocol 4095: If SDW is configured to use HTTPS protocol When installing SMHV on a cluster, the same port number should be used across all nodes.

4.4 ARCHITECTURE

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.

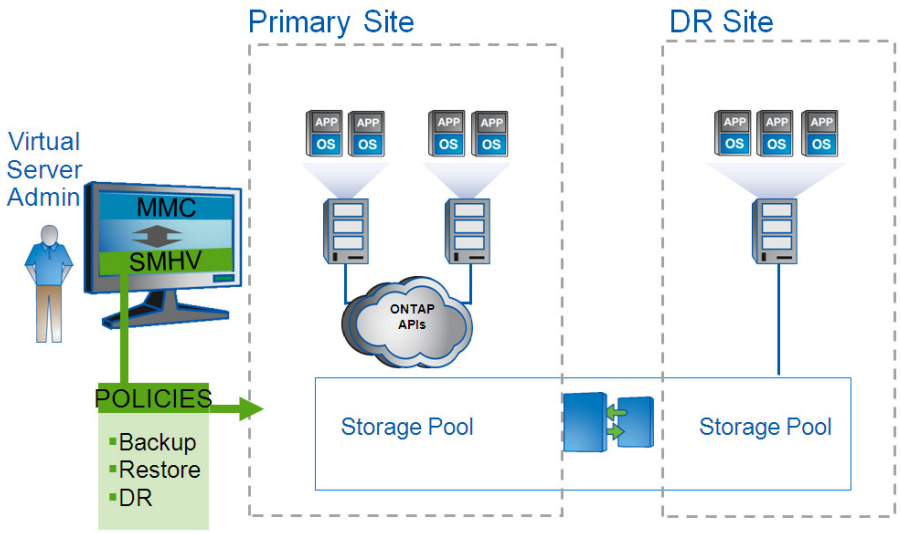

Figure 1) SMHV architecture

4.4.1 COMPONENTS

LICENSE REQUIREMENTS

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

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

NETAPP DATA ONTAP

SMHV will only function within a NetApp storage environment. SMHV requires that the primary storage, where the VMs actually reside, and the secondary storage used as the SnapMirror destination run the Data ONTAP® storage software.

Table 1) Licensing and ONTAP versions

If You Use	Then Use Data ONTAP Version
Host-based licensing	7.3.1P1 or later
Storage system licensing	7.3.2 or later
Storage system licensing with vFiler™	7.3.1.1 P8, 7.3.2P1, or later

For the most current information, see the NetApp Interoperability Matrix Tool (IMT) at http://now.netapp.com/NOW/products/interoperability.

In addition, the following licenses are required:

- SnapRestore
- The required protocol license (FCP, iSCSI)
- SnapMirror (if required)
- SnapDrive 6.2 for Windows (needs to be licensed on the Hyper-V host)

SNAPMANAGER FOR HYPER-V 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): Core and Full
- Windows Server 2008 x64 Standard and Enterprise Editions (Full and Core)
- Windows Server 2008 x64 Standard and Enterprise Editions 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

For the most current information, see the NetApp IMT at http://now.netapp.com/NOW/products/interoperability.

SNAPMANAGER FOR HYPER-V SNAPINFO SETTINGS

SMHV SnapInfo folder stores backup metadata. This can be set up by specifying the SnapInfo settings in the Hosts Management wizard. The metadata information is critical to recovering VMs should a failure occur. 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.

Information
The SnapInfo path must reside on a Data ONTAP LUN. For managing dedicated VMs, the SnapInfo location needs to be a dedicated Data ONTAP LUN. For managing shared VMs, the SnapInfo location needs to be to a shared Data ONTAP LUN. The SnapInfo path must not reside on a CSV.

Information
If SnapInfo settings, are changed you need to manually move all files from the original SnapInfo location to the new location. SnapManager for Hyper-V does not move them automatically.

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

SNAPMANAGER FOR HYPER-V 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.

Information
The report path must not reside on a CSV.

.

SNAPMANAGER FOR HYPER-V EVENT NOTIFICATIONS

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

5 THE SMHV PROCESS FLOW

5.1 ADDING A HYPER-V PARENT HOST OR HOST CLUSTER

If you add a single host, SMHV manages the dedicated VMs on that host. If you add a host cluster, SMHV manages the shared VMs on the host cluster. If you 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, you can configure them after you add the host, using the Configuration wizard. You must configure the backup repository and report directory settings to add and manage VMs using SMHV. Notification settings are optional.

Information
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 you should manage a host from only one management console, if you need to do so from multiple consoles, you can import and export host and dataset configuration information from one remote management console to another to make sure of data consistency. You can also use the Import and Export wizard to change host and dataset configuration settings to a previously exported setting. If you perform this operation in a clustered environment, you need to import the settings on all nodes in the cluster so that all host and dataset configurations are the same. You should not import or export configuration information to the directory where SMHV is installed. If you uninstall SMHV, this file will be lost.

5.2 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 upon 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 is possible. SMHV supports backup of dedicated or clustered VMs. It also supports backups of shared VMs running on CSVs.

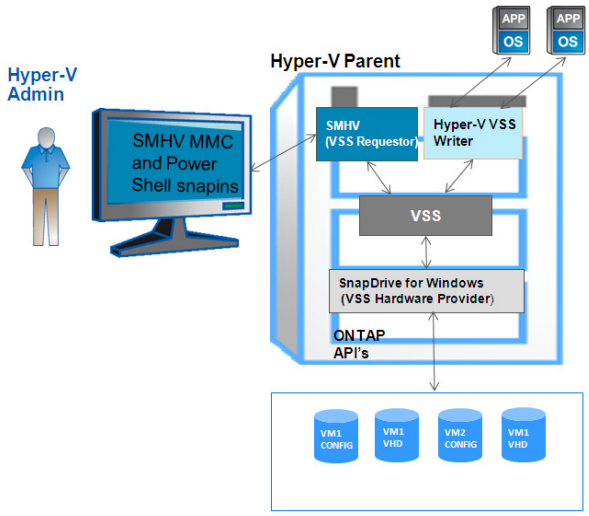

Figure 2) Hyper-V infrastructure and associated storage during an SMHV backup.

Figure 2 represents a high-level overview of the typical SMHV architecture on the primary site storage and will be used in detailing the backup process flow.

1. The SMHV Service is a VSS Requestor which 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. VSS framework requests the hardware provider to mount the LUNs from the Snapshot copy.
5. Hyper-V writer recovers data on the LUNs and brings it to the state of the software Snapshot copy which 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. Upon completion of the local backup, SMHV will update an existing SnapMirror relationship on the volume if the SnapMirror option was selected. SnapMirror will be discussed in further detail in a later section of this document.

SMHV enables you to create application-consistent backups of a VM, if you have Microsoft Exchange, Microsoft SQL, or any other VSS-aware application running on VHDs in the VM. SMHV coordinates with the application VSS writers inside the VM to make sure that application data is consistent when the backup occurs.

Information

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

Information

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

Information

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 gives this option. As part of copying the physical disk contents, it is also copying the disk signature, and this causes the disk signature conflict during the backup. More information is available here:
http://support.microsoft.com/kb/975695
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.
SnapManager for Hyper-V does not support the backup and restore of virtual machines running on SAN boot LUNs. This is a limitation of SDW

Best Practice

When creating a dataset, you should select all VMs that reside on a particular Data ONTAP LUN. This enables you 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 you adds VMs on the same CSV in different datasets, you need to ensure that the backup schedules of these datasets do not overlap

Best Practice

If you change a VM Snapshot copy location to a different Data ONTAP LUN after creating the VM, you should 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.

5.3 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 255 NetApp Snapshot copies per volume limit must be taken into consideration when scheduling backups and configuring the associated retention policies. The number of Snapshot copies per volume can be managed with the proper scheduling and retention policies on a per scheduled backup basis while still meeting SLAs on the VMs.

5.3.1 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 which apply to all hosts that are dataset members.

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

5.3.2 RETENTION POLICIES

The following list describes the retention tags available in SMHV:

Hourly Hourly intervals
Daily A specified time within a 24-hour period
Weekly A specified day and time within a seven-day period
Monthly A specified day and time within a calendar month
Unlimited Backups are never deleted

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

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

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

Information
The option, unlimited, should be used with caution. When selecting this option, backups and the associated NetApp Snapshot copies will be maintained until manually deleted by the administrator. These Snapshot copies are included in the maximum number supported on a volume. Of further note, the NetApp Snapshot copies associated with on demand backups must also be considered when determining the number of Snapshot copies maintained against a volume.

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

For stand-alone Hyper-V hosts:
Registry key: HKLM\SOFTWARE\NetApp\SnapManager for Hyper-V\Server DWORD value: snapinfo_snaps_count (number of SnapInfo Snapshot copies to be retained)

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

5.4 HANDLING SAVED STATE BACKUP OF VMS

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

Information

You can also select the Allow saved state VM backup check box to allow SMHV to back up the VM using the saved state. If you check this option, SMHV will not fail the backup when the Hyper-V VSS writer backs up the VM using the saved state or performs an offline backup of the VM. Doing a saved state or offline backup can cause downtime. For more information on online of offline VM backups, see the Hyper-V Planning for the Backup information in the Technet library:

http://technet.microsoft.com/en-us/library/cc753637(WS.10).aspx

Best Practice

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

5.5 BACKUP SCRIPTS

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

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

Information

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

5.6 QUICK/LIVE MIGRATION IMPLICATIONS

Best Practice

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

5.7 RESTORE PROCESS

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

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

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

1. SMHV restores a VM in coordination with Hyper-V VSS writer. Hyper-V VSS writer will power off the VM and delete it before restore.
2. Files are restored as described above based on restore type.
3. SMHV will notify that the files of the VM are restored properly. Hyper-V VSS writer will the register the VM and VM gets added back in the Hyper-V manager.
4. SMHV starts the VM after restore and executes a post script if specified in the restore wizard.

Information

While restoring the following warning messages maybe displayed:

1. VM to be restored is not [currently running] on the host.
2. VM to be restored is currently running on the host, and:

- It has more VHDs associated to it than at the time of backup.
- It has fewer VHDs associated to 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 path have changed.

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

Information

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

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

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

Best Practice

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

Information

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

5.8 MOUNTING A BACKUP

Backups can be mounted using SnapDrive 6.2 for Windows. The mounted backup is a clone of the protected VM. Once mounted, the backup is displayed within the explorer of Hyper-V host and can be browsed.

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

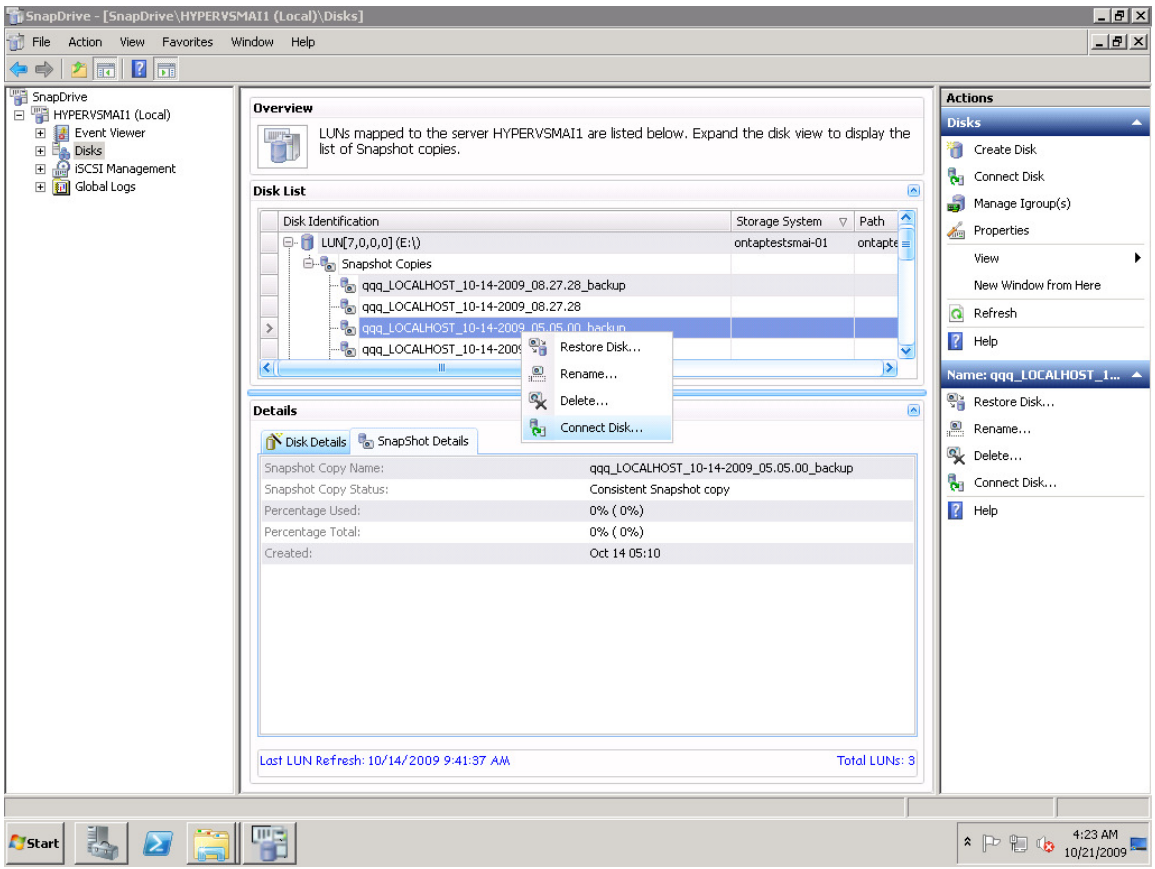

Figure 3) Mounting a backup using SDW.

2. Right-click the Snapshot copy (the one with _backup suffix) and select the connect disk option.

3. In the Provide a Storage System Name, LUN Path and Name panel, perform the following actions:
   a. In the “Storage System Name” field, type the name of the storage system where the LUN will be connected, or choose a storage system from the drop-down list.
   b. In the “LUN Path” field, type the path to the LUN. Alternatively, click Browse and navigate to the LUN you want to connect.
   c. Click Next.

4. If the LUN is a dedicated disk, go to the next step; otherwise, if the LUN is a Windows cluster resource, perform the following steps in the Specify Microsoft Cluster Services Group panel.
   In the Specify Microsoft Cluster Services Group panel, perform ONE of the following actions and then click Next.
   a. Select a cluster group from the Group Name drop-down list.
   b. Select Create a new cluster group to create a new cluster group.
   Note: When selecting a cluster group for your LUNs, choose the cluster group your application will use. If you are creating a volume mountpoint, the cluster group is already selected. This is because the cluster group owns your root volume physical disk cluster resources. It is recommended that you create new shared LUNs outside of the cluster group.
   c. Select Add to cluster shared volumes.

5. In the Select LUN Properties panel, perform the following actions: Either select a drive from the list of available drive letters, or enter a mountpoint for the LUN you are connecting. When you create a volume mountpoint, enter the drive path that the mounted drive will use: for example, G:\mount_drive1\.

6. In the Select Initiators panel, choose an initiator for the LUN.

7. In the Select Initiator Group Management panel, specify whether you will use automatic or manual igroup management.

8. In the Completing the Connect Disk Wizard panel, perform the following actions:
   a. Verify all the settings.
   b. If you need to change any settings, click Back to go back to the previous wizard panels.
   c. Click Finish.

9. Browse the backup by selecting the drive letter on the explorer of Hyper-V host.

SINGLE FILE RESTORE CAPABILITY

In addition to backup verification, mounting a backup provides a way to restore a single file from within a VM on a case-by-case basis. This is performed by attaching a VHD from within the mounted backup as an existing hard drive to a VM within Hyper-V manager.

Once a backup has been mounted, the following manual steps can be followed to perform the restore of a single file, or files, from within a VM backup.

1. Right-click the VM to which the VHD will be attached.
2. Click Settings within the drop-down menu.
3. Select IDE controller option.
4. When asked to choose the type of device you want to add, select Hard Drive and click Add.
5. Select Virtual Hard Disk (.vhd) file and click Browse.
6. Browse to the mounted backup.
7. Select the required VHD file and click OK.
8. The VHD file will be attached to the specified VM after clicking OK.

NetApp recommends creating a “landing” VM that can be used as a target system when mounting backups and attaching preexisting hard drives (VHDs). One VM should be created as a target for each guest operating system within the Hyper-V environment. Once the VHD is mounted, files can be copied across the network back to their desired location. By following this practice, the landing VM can be powered off once the restore is completed without interrupting production, and the VHD removed before unmounting the backup.

Not only is this desirable as it limits disruption to a production VM, but it also prevents a backup from remaining in a mounted state for an extended period of time.

Information
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.

The required steps for removing a VHD are as follows:

1. Power down the VM.
2. Right-click the VM with Hyper-V Manager and select Settings from the drop-down menu.
3. Within the IDE controller window, select the hard disk that needs to be removed. Verify you have selected the correct disk by checking the disk file name before clicking Remove.
4. Click OK.

At this point, the backup can be unmounted using the steps below:

1. Select the disk that you had mounted.
2. Right-click the disk and select the disconnect option.

6 HIGH AVAILABILITY

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

It should be noted that the SMHV, being a host-end application, offers services provided that the storage is continuously available.

Following is the detailed description on the available tools that facilitate storage availability.

6.1 MULTIPATH HA WITH ACTIVE-ACTIVE NETAPP CONTROLLERS

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

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

6.2 DATA ONTAP DSM FOR WINDOWS MPIO

Microsoft MPIO is a protocol-independent feature that supports multiple data paths to a storage device with iSCSI, Fibre Channel, or SAS. Providing multiple paths that can handle failover increases the availability from a host to the storage system. Windows 2008 R2 x 64 servers include support for Microsoft MPIO.

NetApp Data ONTAP device-specific modules (DSMs) for Windows MPIO help NetApp storage systems to integrate with Microsoft MPIO on Windows 2008 R2 server and provides high availability to applications using path failover methods. It determines all the paths pointing to the same LUN so that MPIO can group them into the virtual disk that Windows Server 2008 Hyper-V server will mount. It is also responsible for communicating with MPIO to identify which path 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 Practice
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.

7 APPLICATION CONSISTENCY

Microsoft's Volume Shadow Copy Service, or VSS, was written specifically to enable third-party backup and recovery solutions to provide application-consistent backup and recovery for mission-critical Microsoft supported applications. When VSS is properly configured within the Hyper-V environment, an SMHV initiated Snapshot copy will begin 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 the 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 that is 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, and 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, and 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 allows you to create application-consistent backups of a VM if you have Microsoft Exchange, Microsoft SQL, or any other VSS aware application running on VHDs in the VM. The applications that exist in the VM restore to the same state as 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 will not create 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 you would need to configure application backup products in the VM (for example, SnapManager for Exchange, SnapManager for SQL, and so on).

Information
The Data ONTAP VSS hardware provider is installed automatically as part of the SnapDrive software installation. To make sure the Data ONTAP VSS hardware provider works properly, do not use the VSS software provider on Data ONTAP LUNs. If you use the VSS software provider to create Snapshot copies on a Data ONTAP LUN, you will be unable to delete that LUN using the VSS hardware provider.

.

Information
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 due to a transient problem. If

this event is logged for a failed backup, retry the backup.

Information

SMHV coordinates with Hyper-V VSS writer to create application-consistent backup of VMs. Hyper-V writer communicates with integration services (Hyper-V Volume Shadow Copy 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 storage system.
For details on VM backup, refer to the following TechNet link:
http://technet.microsoft.com/en-us/library/dd252619(WS.10).aspx.

8 CONCLUSION

SnapManager 1.0 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. For more information about any of the solutions or products covered in this report, contact NetApp.

9 SUMMARY OF SNAPMANAGER FOR HYPER-V BEST PRACTICES

Best Practice

The recommendation while deploying Hyper-V on a shared storage is to have one VM per LUN configured. All the VHDs relative to a single VM (VM with multiple drives) can reside on single LUN provisioned as a shared storage to a WFC. It is a best practice for Windows 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 HTTP protocol
4095: If SnapDrive is configured to use HTTPS protocol
The default port number is 808. When installing SMHV 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, you should select all VMs that reside on a particular Data ONTAP LUN. This enables you to get all backups in one Snapshot copy and to reduce the space consumption on the storage system.

Best Practice

If you change a VM Snapshot copy file location to a different Data ONTAP LUN after creating the VM, you should create at least one VM Snapshot copy using Hyper-V manager before creating a backup using SMHV. If you change the Snapshot file location and do not create a VM Snapshot copy before creating a backup, the backup could fail.

Best Practice

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

Best Practice

Choose a backup retention level based on your backup creation and verification schedule. If a Snapshot copy deletion occurs, you should make sure that a minimum of one verified backup remains on the volume. Otherwise, you run a higher risk of not having a usable backup set to restore from 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. Should a backup run against a dataset that has VMs actively being migrated, an error will be generated, and those particular VMs will not be backed up. NetApp recommends that VMs are 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 same, the restored VM might have extra/fewer VHDs. If that is the case NetApp recommends that the cluster configuration of the VM and its dependencies be manually updated.

Best Practice

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

Best Practice

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.

APPENDIX A: VIRTUAL MACHINE MANAGING ITSELF

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

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

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

APPENDIX B: DATA ONTAP VSS HARDWARE PROVIDER REQUIREMENT

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

Viewing Installed VSS Providers

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

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:

vssadmin list providers
The output should be similar to the following:
Provider name: 'Data ONTAP VSS Hardware Provider'
Provider type: Hardware
Provider Id: {ddd3d232-a96f-4ac5-8f7b-250fd91fd102}
Version: 6.2.0.xxxx

Verifying That 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 step.
Navigate to System Tools > Event Viewer > Application in MMC and look for an event with the following values.

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

Information
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 due to a transient problem. If this event is logged for a failed backup, retry the backup.

APPENDIX C: VIRTUAL MACHINE BACKUPS TAKE TOO LONG TO COMPLETE

If a virtual machine contains several direct-attached iSCSI LUNs or pass-through LUNs, and SnapDrive for Windows is installed on the virtual machine, the virtual machine backup can take a long time. The Hyper-V writer takes a hardware snapshot of all the LUNs in the virtual machine using the SnapDrive for Windows VSS hardware provider. There is a Microsoft hotfix that uses the default system provider (software provider) in the virtual machine to make the snapshot. 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 Knowledge Base article 975354 on the Microsoft support site.

Knowledge Base article 975354 - http://support.microsoft.com/