



Technical Report

SnapDrive for Windows Best Practice Guide for Clustered Data ONTAP used in NAS Environments

Vinith Menon, NetApp
January 2015 | TR-4218

TABLE OF CONTENTS

1	Overview	4
2	Purpose and Scope	4
3	SnapDrive for Windows Key Features	4
4	Clustered Data ONTAP	4
5	New in Clustered Data ONTAP 8.3	5
6	Features in Clustered Data ONTAP	5
7	Clustered Data ONTAP and CIFS SVM Setup and Configuration	9
7.1	General Considerations of Clustered Data ONTAP for SMB 3.0 Workloads	9
7.2	Root and Data Volume Settings	10
7.3	Data and Management LIF Settings	10
7.4	SMB 3.0 Settings	10
7.5	Offloaded Data Transfer Settings	11
7.6	Remote Volume Shadow Copy Service Settings (Shadow Copy Feature VSS)	11
7.7	Automatic Node Referral Settings	12
7.8	Creating Network Interface Failover Groups	12
8	New in SnapDrive 7.1 for Windows	13
8.1	Clone of Clone	13
8.2	Split Clone	14
9	SnapDrive for Windows Architecture	16
9.1	Prerequisites for Installation of SnapDrive for Windows	16
9.2	Supported Configurations	17
9.3	Architecture of SnapDrive for Windows	17
10	SnapDrive for Windows and Windows PowerShell	18
10.1	Provisioning Templates	19
11	Adding a Storage Virtual Machine in SnapDrive	21
12	Provisioning a Volume and a CIFS Share Using SnapDrive for Windows	21
13	Backup and Restore Architecture of a CIFS Share Using SnapDrive for Windows	22
13.1	Backup and Restore Architecture	22
13.2	Backup and Restore	22
13.3	Mounting and Dismounting a Snapshot Copy	23
14	Creating VMs in CIFS Shares Created by SnapDrive for Windows	24

15 SnapMirror and SnapVault	25
15.1 SnapMirror	25
15.2 SnapVault	25
16 Restoring Snapshot Copies After Storage Live Migration	26
Version History	26

LIST OF TABLES

Table 1) Supported configurations for SAN and SMB workflows with SnapDrive 7.1 for Windows.	17
Table 2) PowerShell cmdlets.....	18

LIST OF FIGURES

Figure 1) Offloaded data transfer.	8
Figure 2) Each storage virtual machine has its own QoS policy.....	9
Figure 3) Clone-of-clone and split-clone features.....	14
Figure 4) Architecture of SnapDrive for Windows.....	18
Figure 5) VHDX location settings.	24
Figure 6) VM location settings.	24
Figure 7) SnapVault workflow.....	25

1 Overview

NetApp® SnapDrive® for Windows® (SDW) helps you to perform storage-provisioning tasks and manage data in Microsoft® Windows environments. You can run SnapDrive software on Windows hosts in either a physical or a virtual environment. SnapDrive software integrates with Windows Volume Manager so that storage systems can serve as virtual storage devices for application data in Windows Server® 2008 R2, Windows Server 2012, and Windows Server 2012 R2. You can also use Snapdrive to provision storage for Windows virtual machines hosted on ESX® hypervisors.

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

In addition to provisioning in SAN environments, SnapDrive also provisions Server Message Block (SMB) 3.0 shares to support the provisioning and protection of Hyper-V® over SMB and SQL Server® over SMB workloads.

2 Purpose and Scope

SDW performs storage-management tasks and enables application-consistent backup and restores by integrating with NetApp SnapManager® products. It also enables the replication of NetApp Snapshot® copies to remote storage for both SAN and SMB environments. The purpose of this document is to enable users who use SnapDrive 7.x for Windows to deploy workloads such as Hyper-V over SMB, CIFS shares, and SQL Server® over SMB.

3 SnapDrive for Windows Key Features

SnapDrive for Windows provides the following features in clustered NetApp Data ONTAP® environments:

- Online storage configuration, logical unit number (LUN) expansion, and streamlined management
- Integration of Data ONTAP Snapshot technology with the Microsoft Volume Shadow Copy Service (VSS) framework and thus creation of point-in-time images of data stored on LUNs
- Integration with the Microsoft remote VSS framework to perform backup and restore of SMB 3.0 shares hosted on NetApp storage systems running clustered Data ONTAP 8.2 and later
- Support for Microsoft cluster configurations
- Support for iSCSI, FC, and FCoE in Windows and VMware® environments
- Support for virtual Fibre Channel adapters for guest virtual machines in Windows Server 2012 and Windows Server 2012 R2 environments
- Support for raw device mapping (RDM), virtual machine disk (VMDK) over the Network File System (NFS), and Virtual Machine File System (VMFS) datastores in VMware environments
- Native NetApp SnapVault® integration
- Group-managed service accounts in Windows Server 2012
- IPv6 support
- Data ONTAP PowerShell™ cmdlets to run SAN and SMB workflows

4 Clustered Data ONTAP

This section introduces the clustered Data ONTAP architecture and some key concepts for clustered Data ONTAP in the context of SnapDrive.

Storage controllers running clustered Data ONTAP are referred to as nodes. Nodes are joined together in a clustered system; the nodes communicate continuously with each other, coordinate cluster activities, and transparently move data between nodes. Although the basic unit of a cluster is the node, nodes are added to the cluster as part of a high-availability (HA) pair.

As with Data ONTAP operating in 7-Mode, HA pairs enable high availability by communicating with each other over an HA interconnect (separate from the dedicated cluster network) and by maintaining redundant connections to the HA pair's disks. Also, as with Data ONTAP operating in 7-Mode, disks are not shared between HA pairs, although shelves may contain disks that belong to either member of a pair.

Clusters are administered on a whole-cluster rather than a per-node basis, and data is served from one or more storage virtual machines (SVMs). Each SVM is configured to own storage, in the form of volumes provisioned from a physical aggregate and logical interfaces (LIFs) assigned either to a physical Ethernet network or to Fibre Channel target ports. LUNs are created inside an SVM's volumes and are mapped to hosts to provide them with storage space. SVMs are node independent and cluster based. Therefore, they can make use of physical resources, such as volumes or network ports, anywhere in the cluster.

SnapDrive 7.x for Windows supports clustered Data ONTAP 8.2 and later.

For more information about best practices for setting up clustered Data ONTAP, refer to [TR-3450: High-Availability Pair Controller Configuration Overview and Best Practices](#).

5 New in Clustered Data ONTAP 8.3

The recent introduction of NetApp MetroCluster™ support for clustered Data ONTAP 8.3 provides a highly efficient and reliable disaster recovery solution. It works across applications; in virtual environments; and across storage types, protocols, and tiers. MetroCluster provides continuous availability for critical applications for which either planned or unplanned downtime is unacceptable.

The minimum configuration for MetroCluster is a disaster recovery group containing one high-availability (HA) pair at each of two sites, for a total of four nodes (controllers). Each cluster is an active-active HA pair, so all nodes serve clients at all times. This solution can stretch across data centers deployed up to a maximum distance of 124 miles (200 km). This reach enables a level of availability beyond the capabilities of a local cluster, making MetroCluster a highly versatile solution.

Because MetroCluster is an active-active solution, all nodes in each cluster actively serve data to applications, and data can be read from both the primary and secondary clusters, a feature that can also improve read performance.

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

6 Features in Clustered Data ONTAP

One of the major components added to clustered Data ONTAP 8.2 and later is support for the SMB 3.0 NAS protocol. This support enables NetApp customers to use the SMB 3.0 features introduced with Microsoft Windows Server 2012 and Windows Server 2012 R2. With these new features, you can use clustered Data ONTAP to host VM virtual disks and configuration settings on a CIFS file share.

Some of the SMB 3.0 features implemented in clustered Data ONTAP 8.2 and later that support continuously available file shares and Hyper-V storage include the following:

- Persistent handles (continuously available file shares)
- The witness protocol
- Clustered client failover
- Scale-out awareness
- Remote VSS

Hyper-V over SMB and SQL Server over SMB workloads are supported only with the following configuration:

- SMB 3.0
- Clustered Data ONTAP 8.2 and later
- Windows Server 2012 and Windows Server 2012 R2

Persistent Handles (Continuously Available File Shares)

To enable continuous availability on a file share, the SMB client opens a file on behalf of an application, such as a VM running on a Hyper-V host, and requests persistent handles for the virtual hard disk format (VHDX) file. When the SMB server receives a request to open a file with a persistent handle, the SMB server retains sufficient information about the file handle, along with a unique resume key supplied by the SMB client. Persistent-handle information is shared between nodes in a cluster.

In the case of a planned move of file share resources from one node to another or the failure of a node, the SMB client reconnects to an active and available node and uses persistent handles to reopen the file. The application or VM running on the SMB client computer does not experience any failures or errors during this operation. From the perspective of the VM, I/O operations to virtual disks appear to be delayed for a short period of time in a way that is similar to a brief loss of connectivity to the disk, but no disruption is noticed.

Note: Continuously available shares are not supported for SQL Server over SMB environments.

Witness Protocol

When an SMB server node fails, the SMB client usually relies on the Transmission Control Protocol (TCP) timeout to detect a failure of the file share resource such as an open file. SMB 3.0 allows variable values for TCP timeouts, and, because the virtual disk is a critical resource, a VM running on a Hyper-V server requires faster detection of network-resource failover. In addition, the witness protocol significantly improves the SMB client reconnect time.

During connection to a shared resource (for example, by using `TREE_CONNECT`), the SMB server provides information about features enabled on a share, such as whether the resource is clustered, scaled out, or continuously available. Based on this information, the SMB client requests this same data from the other nodes. Upon receiving the information, the SMB client registers itself with the other nodes.

In the case of a cluster node failure, the SMB client is already connected to another node, which can detect the failure and then notify the SMB client. This saves the SMB client from waiting until the TCP timeout is over and instead initiates an immediate reconnect to the running node, minimizing the time the client is disconnected from the resource. For VMs with virtual disks stored on such SMB shares, the disk disconnection time is reduced to the point that the VMs do not detect such disconnects as hardware failure.

This feature is enabled on clustered Data ONTAP by default only if all best practices are followed and there is a logical interface (LIF) on each node in the cluster in every storage virtual machine (SVM; formerly called Vserver). In addition, the witness protocol comes into play only for continuously available shares.

Clustered Client Failover

To increase redundancy in a VM environment, place Hyper-V servers into a Microsoft failover cluster. When the Hyper-V server node running a VM fails, the VM is live migrated to another node. Prior to the introduction of clustered client failover (CCF) with SMB 3.0, a VM moving to another cluster node was considered to be a new application instance. New application instances connecting to files already opened on file shares had to wait until the TCP timeout was over and the file handle was closed. CCF allows the VM to open a virtual disk file on a file share and provides a unique application identifier. When

a Hyper-V server cluster node fails, the VM starts on another Hyper-V server node and supplies the same application identifier, letting the SMB server close existing file handles. The SMB client can then reconnect to the previously open file.

Scale-Out Awareness

Clustered Data ONTAP supports scale-out and can serve data from multiple nodes. It brings additional data redundancy to the network and spreads the load of multiple SMB clients between multiple nodes in a cluster. Scale-out awareness allows SMB clients to connect to all nodes in the cluster and access the same data.

Remote Volume Shadow Copy Service

VSS is a framework that coordinates application I/O and physical storage on the same server and allows the creation of application-consistent Snapshot copies of the storage. Windows Server 2012 and Windows Server 2012 R2 extend the functionality of VSS to multiple servers. For example, an application running on one server has storage on another server's file share. Remote VSS coordinates I/O activities during a backup process between both servers and provides application-consistent Snapshot copies of the storage for applications running remotely on the storage server. Clustered Data ONTAP 8.2 and later extend the functionality of remote VSS by plugging into the VSS framework. A VSS service runs on a NetApp controller, and a VSS provider runs on Windows Server 2012 and Windows Server 2012 R2 machines. From a VSS perspective, the NetApp array acts exactly the same as a Windows file server. SDW uses remote VSS to back up and restore workloads hosted on SMB shares.

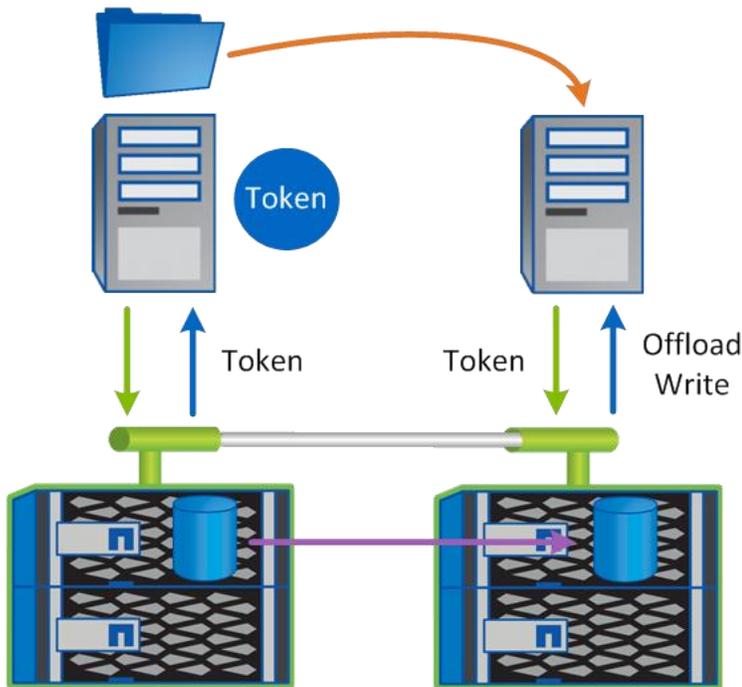
Offload Data Transfer

Copy offload provides a mechanism for performing full-file or subfile copies between two directories residing on remote servers; the servers can be the same or different. Here, the copy is created by copying data between the servers or within the same server if both source and destination files are on the same server, without the client reading the data from the source and writing to the destination. This strategy reduces utilization of client-server processor and memory utilization and minimizes network I/O bandwidth.

With Windows Server 2012 and Windows Server 2012 R2, Microsoft introduced a copy-offload mechanism capable of making copies between two independent servers, provided there is an underlying mechanism to move data between the servers. The previous server-side copy mechanism required both the source and destination file to be on the same server.

Before proceeding with a copy operation on the host, make sure that copy-offload settings are configured on the storage system. SnapDrive or Windows does not enable copy-offload settings on the storage system directly. After provisioning LUN shares using SnapDrive, activities such as storage live migration and file copy are offloaded to the storage system, and SnapDrive does not participate in these operations. Figure 1) Offloaded data transfer. depicts the process of offloaded data transfer.

Figure 1) Offloaded data transfer.



Storage Quality of Service

Storage quality of service (QoS) in clustered Data ONTAP 8.2 (and later) enables the grouping of storage objects and the setting of throughput limits on the group. With this capability, a storage administrator can separate workloads by organization, application, or business unit. An administrator can also distinguish between production and development environments.

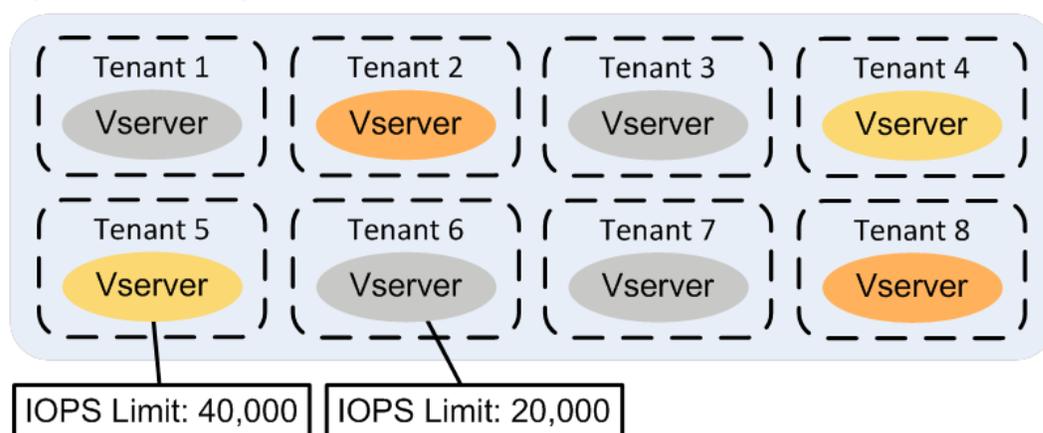
In enterprise environments, storage QoS helps to:

- Prevent user workloads from affecting each other.
- Protect critical applications that have specific response times that must be met in IT-as-a-service (ITaaS) environments.
- Prevent tenants from affecting each other.
- Avoid performance degradation with each new tenant.

QoS allows you to limit the amount of I/Os sent to an SVM, a flexible volume, a LUN, or a file. I/Os can be limited by the number of operations or by the raw throughput. Each SVM can have its own QoS policy, as indicated in Figure 2.

After the user configures SVMs or volumes for QoS, SDW 7.x can provision storage from these SVMs onto the host system or the guest VMs after providing details for these SVMs to the host. For QoS best practices for Microsoft applications, refer to the appropriate [SnapManager Best Practice Guides](#).

Figure 2) Each storage virtual machine has its own QoS policy.



IPv6 Support

Clustered Data ONTAP 8.2 and later support IPv6. SnapDrive 7.x for Windows supports IPv6 in all workflows that require the addition of an IP address or host name; examples include the addition of a storage system in the transport protocol settings. Use the following command to enable IPv6 in a clustered Data ONTAP 8.2 (and later) cluster.

```
network options ipv6 modify -enabled true
```

Verify that the DNS server has either the IPv6 or IPv4 address of the host system, but not both.

Note: SnapDrive does not support mixed-mode IP formats. This means that both the host and storage system must use the same IP format (IPv4 or IPv6).

Note: After IPv6 has been enabled on the storage cluster, it cannot be disabled.

Note: In IPv6 environments, NetApp recommends providing a fully qualified domain name in SnapDrive.

For more information, refer to the [Clustered Data ONTAP Network Management Guide](#).

7 Clustered Data ONTAP and CIFS SVM Setup and Configuration

There are no special requirements for a clustered Data ONTAP 8.2 and later setup to use SMB 3.0 features. By default, clustered Data ONTAP 8.2 and later support all versions of SMB. Most of the usual applications for the SMB protocol, such as user file sharing, can work on SMB protocols earlier than 3.0, and these applications might not benefit from the additional features provided by SMB 3.0. Hyper-V workloads require the SMB 3.0 protocol and some of its additional features, such as continuously available shares. Considering the additional overhead of storing and replicating persistent-handle information between nodes in an HA pair to support features such as continuously available shares, NetApp strongly recommends using only continuously available shares for Hyper-V over SMB workloads.

7.1 General Considerations of Clustered Data ONTAP for SMB 3.0 Workloads

When setting up clustered Data ONTAP 8.2 and later for the use of continuously available file shares to host VHDX disk images of VMs, consider the following guidelines.

- Persistent handles work only between nodes in an HA pair.
- The witness protocol works only between nodes in an HA pair.
- Continuously available file shares are supported only for Hyper-V workloads.

- Offloaded data transfer (ODX) is supported with clustered Data ONTAP 8.2 and later and works across protocols. The process of copying data between a file share and iSCSI or an FCP-attached LUN uses ODX.
- NetApp recommends having connectivity between Hyper-V hosts and the NetApp array on a 10GB network, if one is available. In the case of 1GB network connectivity, NetApp recommends creating an interface group consisting of multiple 1GB ports.
- Install CIFS and NetApp FlexClone® technology (required by remote VSS in SnapManager in Hyper-V [SMHV]) licenses.
- Set up time settings on nodes in the cluster. Also, use the Network Time Protocol (NTP) if the NetApp CIFS server has to participate in the Windows Active Directory® domain.

The minimum Microsoft OS versions supporting SMB 3.0 are Windows Server 2012 and Windows Server 2012 R2.

7.2 Root and Data Volume Settings

NetApp CIFS SVM root and data volumes should be NetApp FlexVol® volumes and have an NTFS security style. Symlinks, hardlinks, and widelinks are not supported with NetApp SnapManager for Hyper-V. Also, junctions inside data volumes are not supported.

7.3 Data and Management LIF Settings

Create at least one data LIF per node for every SVM in the cluster. Do not configure the data LIF to AutoRevert. Give each LIF's IP address an entry in the DNS, but do not give it a NetBIOS alias because NetBIOS aliases are not allowed for DNS entries. Network interface failover groups can be configured to specify the network ports to which the LIF can be moved.

7.4 SMB 3.0 Settings

Although the SMB 3.0 protocol is enabled by default, it can be checked, enabled, or disabled by using the `vserver cifs options` command in the advanced mode.

To set up the advanced mode, use the command `set advanced`.

```
Vespus::> set advanced
Warning: These advanced commands are potentially dangerous; use them only when directed to do so
by NetApp personnel.
Do you want to continue? {y|n}: y
Vespus::~*>
```

Advanced mode is required for showing and modifying all other settings.

To verify that SMB 3.0 is enabled, use the command `vserver cifs options show`.

```
Vespus::~*> vserver cifs options show
Vserver: nacifs
Copy Offload Enabled: true
Default Unix Group: -
Default Unix User: pcuser
Export Policies Enabled: false
Is Referral Enabled: false
Is Local Auth Enabled: true
Is Local Users and Groups Enabled: true
Max Multiplex Count: 255
Read Grants Exec: disabled
Shadowcopy Dir Depth: 5
Shadowcopy Enabled: true
SMB2 Enabled: true
```

```
SMB3 Enabled: true
WINS Servers: -
Is Use Junction as Reparse Point Enabled: true
Vespus::*>
```

To enable or disable SMB 3.0, use the command `vserver cifs options modify -vserver <vserver name> -smb3-enabled {true|false}`.

```
Vespus::*> vserver cifs options modify -vserver nacifs -smb3-enabled true
Vespus::*>
```

Note: Continuously available shares are not supported for SQL Server over SMB environments.

7.5 Offloaded Data Transfer Settings

To use ODX for fast provisioning of VMs from the master image prepared with the Microsoft SysPrep utility on the same file share hosting VHDX files, the ODX feature is enabled globally or on a per-SVM basis by default.

To check if ODX is enabled, use the command `vserver cifs options show`.

```
Vespus::*> vserver cifs options show
Vserver: nacifs
    Copy Offload Enabled: true
    Default Unix Group: -
    Default Unix User: pcuser
    Export Policies Enabled: false
    Is Referral Enabled: false
    Is Local Auth Enabled: true
    Is Local Users and Groups Enabled: true
    Max Multiplex Count: 255
    Read Grants Exec: disabled
    Shadowcopy Dir Depth: 5
    Shadowcopy Enabled: true
    SMB2 Enabled: true
    SMB3 Enabled: true
    WINS Servers: -
    Is Use Junction as Reparse Point Enabled: true
Vespus::*>
```

To enable or disable the ODX feature, use the command `vserver cifs options modify -vserver <vserver name> -copy-offload-enabled {true|false}`.

```
Vespus::*> vserver cifs options modify -vserver nacifs -copy-offload-enabled true
Vespus::*>
```

7.6 Remote Volume Shadow Copy Service Settings (Shadow Copy Feature VSS)

Enable this feature to protect VMs if SMHV is deployed on Hyper-V servers.

To check if remote VSS is enabled, use the command `vserver cifs options show`.

```
Vespus::*> vserver cifs options show
Vserver: nacifs
    Copy Offload Enabled: true
    Default Unix Group: -
    Default Unix User: pcuser
    Export Policies Enabled: false
```

```

        Is Referral Enabled: false
        Is Local Auth Enabled: true
    Is Local Users and Groups Enabled: true
        Max Multiplex Count: 255
        Read Grants Exec: disabled
    Shadowcopy Dir Depth: 5
        Shadowcopy Enabled: true
            SMB2 Enabled: true
            SMB3 Enabled: true
            WINS Servers: -
    Is Use Junction as Reparse Point Enabled: true
Vespus::*>

```

To enable or disable the remote VSS feature, use the command `vserver cifs options modify -vserver <vserver name> -shadowcopy-enabled {true|false}`.

```

Vespus::*> vserver cifs options modify -vserver nacifs -shadowcopy-enabled true
Vespus::*>

```

7.7 Automatic Node Referral Settings

The Microsoft Hyper-V host relies heavily on Kerberos authentication that cannot be used with NetApp IP-based automatic node referral. By default, node referrals are disabled, but verify that this is the case when deploying Hyper-V over SMB by using the command `vserver cifs options show`.

```

Vespus::*> vserver cifs options show
Vserver: nacifs
    Copy Offload Enabled: true
    Default Unix Group: -
    Default Unix User: pcuser
    Export Policies Enabled: false
    Is Referral Enabled: false
    Is Local Auth Enabled: true
    Is Local Users and Groups Enabled: true
    Max Multiplex Count: 255
    Read Grants Exec: disabled
    Shadowcopy Dir Depth: 5
    Shadowcopy Enabled: true
        SMB2 Enabled: true
        SMB3 Enabled: true
        WINS Servers: -
    Is Use Junction as Reparse Point Enabled: true
Vespus::*>

```

To disable this feature, use the command `vserver cifs options modify -vserver <vserver name> -is-referral-enabled false`.

```

Vespus::*> vserver cifs options modify -vserver nacifs -is-referral-enabled false
Vespus::*>

```

7.8 Creating Network Interface Failover Groups

To create a network interface failover group, use the `network interface failover-groups` command.

```

Vespus:::> network interface failover-groups create -failover-group nacifs_ela -node Vespus-01 -port ela
Vespus:::> network interface failover-groups create -failover-group nacifs_ela -node Vespus-02 -port ela

```

```

Vespus::> network interface failover-groups show
Failover
Group      Node      Port
-----
clusterwide
  Vespus-01  e0a
  Vespus-01  e0b
  Vespus-01  e0c
  Vespus-01  e0d
  Vespus-01  e1a
  Vespus-02  e0a
  Vespus-02  e0b
  Vespus-02  e0c
  Vespus-02  e0d
  Vespus-02  e1a
nacifs_e1a
  Vespus-01  e1a
  Vespus-02  e1a
12 entries were displayed.

```

8 New in SnapDrive 7.1 for Windows

The following new features have been added to SDW 7.1.

8.1 Clone of Clone

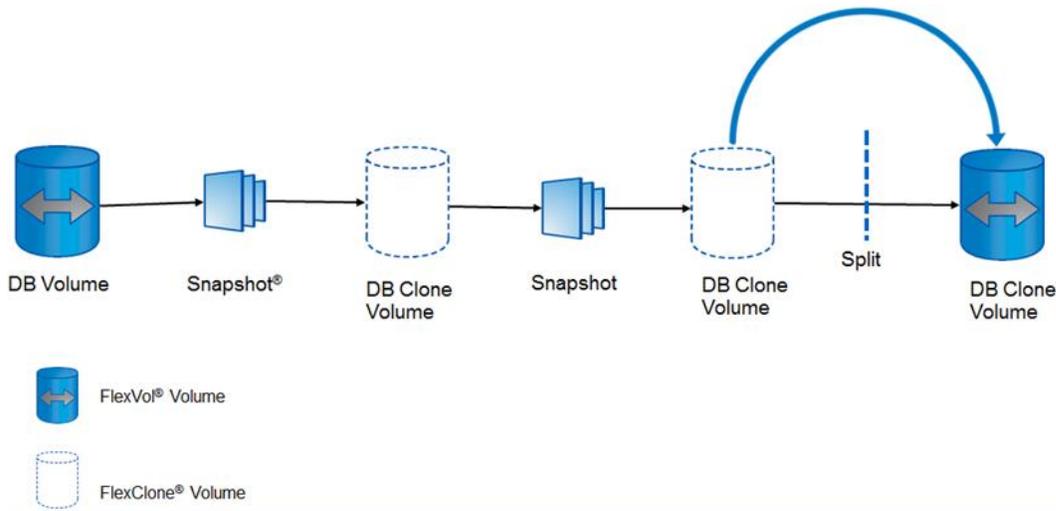
The clone-of-clone feature enables you to provision a volume, create a Snapshot copy of that volume, and then clone the volume by using the Snapshot copy. Data ONTAP supports the clone-of-clone feature. There are two clone types supported on clustered Data ONTAP: Sis-Clone and Volume-Clone (FlexClone). Data ONTAP creates a new volume to host the data from the Snapshot copy from which it was created. This Volume-Clone is tightly bound to the parent volume because the underlying blocks are shared. On top of this Volume-Clone, the user can create another Snapshot copy and mount it as a Volume-Clone. SnapManager for SQL Server uses the clone-of-clone functionality to clone data for testing and development purposes.

Note: Creating a LUN clone from SnapDrive is not supported.

Points to remember:

- By default, SnapDrive only supports two levels of clones. In the case of VMDK on VMFS, however, more than two levels of clones can be created.
- The user must verify that there is adequate space in the aggregate for clone creation.
- SnapDrive doesn't allow creation of a LUN on a Volume-Clone.
- The clone-of-clone operation cannot be performed in a pass-through disk connected to a highly available VM.

Figure 3) Clone-of-clone and split-clone features.



The following sections describe the PowerShell cmdlets that come with SDW 7.1 to support the clone-of-clone feature.

Mount-SdSnapshot Cmdlet

The `Mount-SdSnapshot` cmdlet comes with the `-ValidateCloneDepth` parameter to verify and validate the depth of the clone of clone created.

```
PS C:\Users\administrator.SDSMQA> Mount-SdSnapshot -Path \\CIFS_KIRAN\L1_share111031407110063812 -ValidateCloneDepth -Path \\CIFS_KIRAN\L1_COC_Share2
cmdlet Mount-SdSnapshot at command pipeline position 1
Supply values for the following parameters:
(Type !? for Help.)
Snapshot: COC_L1_Snap1

Resource                               ConnectedResource
-----                               -
\\CIFS_KIRAN\L1_share111031407110063812  \\CIFS_KIRAN\L1_COC_Share211031407124757013
```

Dismount-SdSnapshot Cmdlet

The `Dismount-SdSnapshot` cmdlet comes with the `-DeleteParentClones` parameter to delete the parent clones and dismount the snapshot backup.

```
PS C:\Users\administrator.SDSMQA> Dismount-SdSnapshot -Path \\CIFS_KIRAN\L1_COC_Share211031407124757013 -DeleteParentClones
Dismount-SdSnapshot
'\\CIFS_KIRAN\L1_COC_Share211031407124757013' will be dismounted.
Do you want to continue?
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"):
Dismount-SdSnapshot
'\\CIFS_KIRAN\L1_COC_Share211031407124757013' will be dismounted.
Offline parent volume clones will be deleted, are you sure you want to continue?
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help (default is "Y"):
Successfully dismounted the Snapshot backup.
```

8.2 Split Clone

The split-clone feature enables you to split a Volume-Clone from its parent volume and make it an independent NetApp FlexVol volume. Although SnapDrive provides a command line, there is no GUI for this specific feature. When the Snapshot copy is mounted as a NetApp FlexClone volume, the volume clone depends heavily on the parent volume because they share blocks. The split-clone feature allows

you to remove the dependency of the FlexClone volume by splitting the clone from the parent. All SnapManager products make use of this feature, as does Data ONTAP.

Points to remember:

- During a Volume-Clone split, LUN provisioning and Snapshot management operations are not supported on that volume
- The status of the split clone can be obtained from the command `sdcli clonesplit status`
- Before performing split-clone operations, verify that there is enough space in the aggregate for the operation to succeed
- Because split-clone operations require a significant amount of time (a 1GB volume split-clone operation takes five minutes), be sure to account for this time in any service-level agreements.

The following sections describe the PowerShell cmdlets that come with SDW 7.1 to support the split-clone feature.

Get-SdVolumeCloneSplitEstimate Cmdlet

The `Get-SdVolumeCloneSplitEstimate` cmdlet estimates the space in the aggregate that is needed for the Volume-Clone split operation.

```
PS C:\Users\administrator.SDSMQA> Get-SdVolumeCloneSplitEstimate \\CIFS_KIRAN\L1_share111031406450978811

ResourceName : SDW71_TOI_SMB_VolSiClone78f7f268_4548_4bb7_984f_7bad355cfd57
Message      : Success.
OwnedSpace   :
SharedSpace  :
RequiredSpace : 3725131776
AvailableSpace :
Container    :
```

Start-SdVolumeCloneSplit cmdlet

The `Start-SdVolumeCloneSplit` cmdlet is used to start the Volume-Clone split to create a FlexClone volume.

```
PS C:\Users\administrator.SDSMQA> Start-SdVolumeCloneSplit \\CIFS_KIRAN\L1_share111031406450978811

ResourceName          Message
-----
SDW71_TOI_SMB_VolSiClone78f7f268_4548_4bb7_984f_7bad355c... Success. in_progress,7934;
```

Get-SdVolumeCloneSplit Cmdlet

The `Get-SdVolumeCloneSplit` cmdlet retrieves the status of the Volume-Clone split operation.

```
PS C:\Users\administrator.SDSMQA> Get-SdVolumeCloneSplit \\CIFS_KIRAN\L1_share111031406450978811

ResourceName : SDW71_TOI_SMB_VolSiClone78f7f268_4548_4bb7_984f_7bad355cfd57
PercentComplete : 0
TotalBlocks     : 0
BlocksDone      : 101748
Message         : Success.
JobId           :
```

9 SnapDrive for Windows Architecture

SDW has undergone significant architectural changes from the previous release to provide support for new features, such as Hyper-V over SMB 3.0, and the backup and restore of workloads hosted on CIFS shares by using remote VSS.

SDW can now be used to provision and protect SMB workloads (Hyper-V over SMB and SQL Server over SMB) in addition to its already existing support for SAN systems. The SnapManager suite of products, including SnapManager for Hyper-V and SnapManager for SQL Server, uses SDW to protect SMB-related workloads.

For best practices when setting up Windows Server 2012 Hyper-V over SMB workloads on NetApp storage systems, refer to [TR-4172: Microsoft Hyper-V over SMB 3.0 with Clustered Data ONTAP: Best Practices](#).

For best practices concerning the protection of Hyper-V VMs in SMB environments that use SnapManager for Hyper-V 2.1, refer to [TR-4355: NetApp SnapManager 2.1 for Hyper-V on Clustered Data ONTAP 8.3](#).

For best practices concerning the protection of SQL Server workloads in SMB environments using SnapManager for SQL Server, refer to [TR-4353: Best Practices Guide for Microsoft SQL Server and SnapManager 7.1 for SQL Server with Clustered Data ONTAP](#).

9.1 Prerequisites for Installation of SnapDrive for Windows

.Net 3.5 SP1 and 4.0

.Net 3.5 SP1 and 4.0 are required for installing SDW.

Note: Upgrading .Net might require a reboot depending on the state of the system.

Licenses

Verify that the following licenses are installed on the storage system when performing SMB-related operations using SnapDrive.

- CIFS
- FlexClone (for mount operations and remote VSS)
- NetApp SnapRestore®
- SnapManager suite license on the storage system or a SnapDrive and SnapManager host license on the server
- NetApp SnapMirror® and SnapVault® (optional)

Storage-Side Settings

The process of configuring the storage-side settings involves the following tasks:

1. Configure all parameters as for clustered Data ONTAP.
2. Verify that there is a separate LIF created for SMB communication.
3. Verify that there is at least one data LIF on every node that has a share for nondisruptive operations.

Note: You can simplify the LIF configuration process by using System Manager.

9.2 Supported Configurations

For all SMB 3.0–related workloads, users must use the PowerShell cmdlets or templates that are installed after SnapDrive is installed. The templates are installed at `C:\Program Files\NetApp\SnapDrive\templates`.

PowerShell cmdlets support SnapDrive provisioning; Snapshot copy management, backup, restore, and mounting; and SnapMirror and SnapVault operations for both SAN and SMB. Table 1 indicates the configurations supported for SAN and SMB workflows with SnapDrive 7.1 for Windows.

Table 1) Supported configurations for SAN and SMB workflows with SnapDrive 7.1 for Windows.

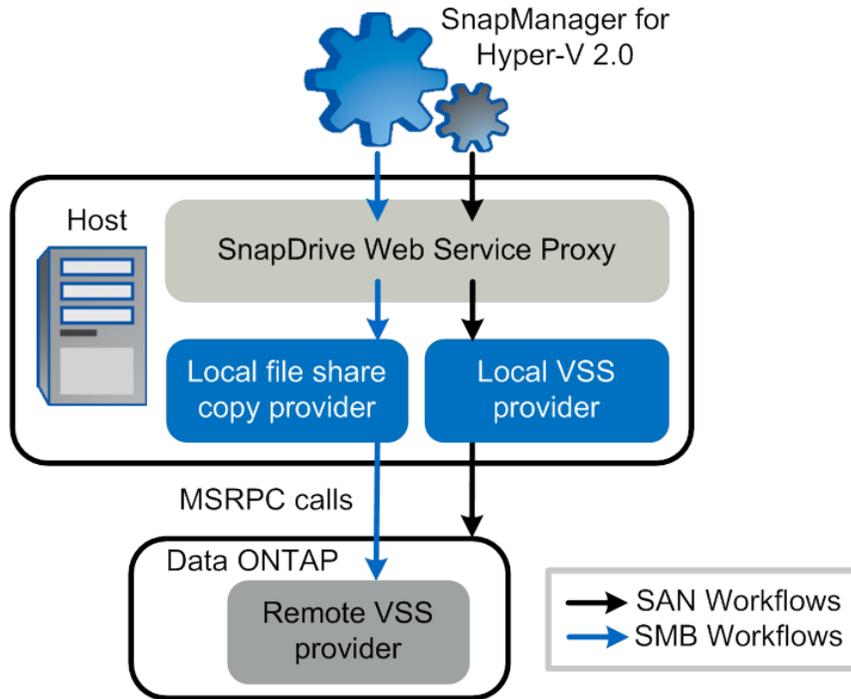
SnapDrive Interface	SAN (FC, iSCSI, and FCOE)	SMB 3.0 (Windows Server 2012 and Windows Server 2012 R2 Only)
SnapDrive GUI	Supported	Not supported
SnapDrive CLI (SDCLI)	Supported	Not supported
SnapDrive PowerShell cmdlets	Basic operations supported	All operations and workflows are fully supported

9.3 Architecture of SnapDrive for Windows

SDW supports both SAN and SMB workflows. As discussed in the previous section, you can access all SMB-related workflows through the PowerShell interface. You can perform all SAN-based workflows by using either the SnapDrive GUI or SDCLI commands. In the case of SMB workflows, the web service proxy passes requests such as backup, restore, and so on to the local file share copy provider that communicates to the remote VSS provider in the Data ONTAP system through Microsoft Remote Procedure Call (MSRPC). If a SAN operation is initiated, these calls are passed to the VSS hardware provider that resides on the host, which then communicates with Data ONTAP for SAN-related operations.

Any SMB or SAN backup, restore, or replication operation initiated through SnapManager products such as SnapManager for Hyper-V or SnapManager for SQL Server reaches the SnapDrive web service proxy layer directly and is redirected accordingly. Figure 4 depicts the architecture of SnapDrive for Windows.

Figure 4) Architecture of SnapDrive for Windows.



10 SnapDrive for Windows and Windows PowerShell

SDW introduces PowerShell cmdlets that can be used to perform common operations and build workflows. This is in addition to the existing SDCLI option. The PowerShell cmdlets listed in Table 2 can be used for both SAN and SMB workflows.

Table 2) PowerShell cmdlets.

PowerShell cmdlets	SAN	NAS	7-Mode System Support
Debug-SdHost	No	Yes	No
Dismount-SdSnapshot	No	Yes	No
Get-SdInfo	Yes	Yes	Yes
Get-SdSMBShadowCopyEmsMessage	No	Yes	No
Get-SdSnapMirror	No	Yes	No
Get-SdSnapshot	Yes	Yes	Yes
Invoke-SdSnapMirrorUpdate	Yes	Yes	Yes, no SnapVault (use Protection Manager)
Invoke-SdEmsAutosupportLog	Yes	Yes	Yes
Get-SdVM	Yes	Yes	Yes
Mount-SdSnapshot	No	Yes	No

PowerShell cmdlets	SAN	NAS	7-Mode System Support
Get-SdStorageConnectionSetting	Yes	Yes	Yes
Get-SdStorage	Yes	Yes	Yes
New-SdSMBShare	No	Yes	No
New-SdSnapshot	Yes	Yes	Yes
New-SdVolume	Yes	Yes	No
Remove-SdSMBShare	No	Yes	No
Remove-SdSnapshot	Yes	Yes	Yes
Remove-SdStorageConnectionSetting	Yes	Yes	Yes
Remove-SdVolume	No	Yes	No
Rename-SdSnapshot	Yes	Yes	Yes
Restore-SdSnapshot	Yes	Yes	Yes
Set-SdSnapshot	Yes	Yes	No
Set-SdStorageConnectionSetting	Yes	Yes	Yes
Get-SdSnapMirrorPolicyRule	Yes	Yes	No
Remove-SdSnapMirrorPolicyRule	Yes	Yes	No
Set-SdSnapMirrorPolicyRule	Yes	Yes	No

For more information about each PowerShell cmdlet, refer to the PowerShell reference guide that is available as part of the SnapDrive software documentation.

Note: Unlike PowerShell cmdlets, SDCLI commands do not support SMB operations or workflows.

10.1 Provisioning Templates

SDW also introduces PowerShell templates that can rapidly provision SMB environments, depending on the workload that the user intends to host on them. The template is configured according to the best practices for the workload deployed on the SMB environment. The PowerShell templates are installed at the following location:

```
C:\Program Files\NetApp\SnapDrive\templates
```

The following templates are available:

- Home directory
- SQL Server over SMB
- Hyper-V over SMB

Note: You can customize these templates depending on the environment, as is shown in the following example:

```
PS C:\Users\administrator.HOST1> New-SdVolume -TemplateName HomeDirProvTemplate.xml -Name test -Size 10GB -Aggregate Aggr1 -StorageSystem 10.1.1.2 -JunctionPath /home_test -VServerContext Vserver1
```

Best Practice

NetApp recommends provisioning all SMB-related workloads, such as CIFS shares for home directories, SQL Server workloads, and Hyper-V workloads, on a host using these PowerShell provisioning templates. This configures the environment according to best practices.

Provisioning Shares for Hyper-V Using New-SdSMBShare cmdlet

This example syntax provisions a share for Hyper-V by using the specified template.

```
New-SdSMBShare -Path / -Name HyperVShare -CIFSServer HyperVFileServer -TemplateName "C:\program
files\SnapDrive\HyperVVHDxProvTemplate.xml"

Acl                : {Everyone / Full Control}
AttributeCacheTtl  : 1
CifsServer         : HyperVFileServer
VServer           : HyperVirtualStorageServer
Comment           : Hyper-V SMB share
DirUmask          : 1
FileUmask         : 1
Path              : /
Volume            : HyperVVolume
ShareName         : HyperVShare
ShareProperties    : {browsable, continuously_available}
SymlinkProperties  : {enable}
UNCPathType       : SMBShare
IsMountedToDrive  : False
MountedDrive      :
ResourceType      : SDSMBShare
ResourceName      : \\HyperVFileServer\HyperVShare
Ranges           :
```

Provisioning a Storage System Volume Using the New-SdVolume cmdlet in the Template

This example provisions a storage system volume using the specified template.

```
New-SdVolume -Name sqldbvolume -Aggregate sqldbaggregate -JunctionPath /sqldbvolume
-TemplateName C:\Program Files\SnapDrive\Templates\HyperVVHDxProvTemplate.xml -Size 128GB -
StorageSystem sqlvirtualstorageserver

Name                : sqldbvolume
Vserver            : sqlvirtualstorageserver
FullPath           : sqlvirtualstorageserver:/vol/sqldbvolume
JunctionPath       : /sqldbvolume
JunctionParentName :
SizeTotal          :
SizeUsed           :
SnapMirrorSource   :
SnapMirrorDest     :
SnapVaultPrimary   :
SnapVaultSecondary :
FlexCloneEnabled   :
IsFlexClone        :
ResourceType       : SDStorageVolume
ResourceName       : sqlvirtualstorageserver:/vol/sqldbvolume
Ranges            :
```

11 Adding a Storage Virtual Machine in SnapDrive

You can execute SMB operations and workflows only after adding the CIFS SVM system to SnapDrive. To do so, run the following PowerShell cmdlet.

```
Set-SdStorageConnectionSetting -Name (Vserver Mgmt LIF or Vserver name)
```

When the CIFS SVM system is added using this cmdlet, a configuration repository file called `Nsf.config` is created at `C:\Program Files\NetApp\SnapDrive\` (the SnapDrive installation folder). This file is subsequently updated when new SVM storage system connections are established or removed.

Note: Verify that the CIFS server name and the SVM name are not identical.

Note: Cluster SVM credentials are not required to be added when adding the CIFS SVM.

Note: NetApp recommends adding the SVM name and IP address in `C:\Windows\System32\Drivers\etc.`

Best Practice

NetApp recommends protecting the `Nsf.config` file by creating another copy of the file and storing it safely. NetApp also recommends updating the backup copy when new storage connections are established or existing ones are removed. After the `Nsf.config` file is restored, SnapDrive services must be restarted.

12 Provisioning a Volume and a CIFS Share Using SnapDrive for Windows

You can provision volumes and shares using standalone PowerShell cmdlets or templates.

Volume Creation and Deletion

Create volumes for hosting SMB shares by using the `New-SDVolume` cmdlet, as follows:

```
New-SdVolume -Name <vol_name> -Aggregate <aggregate_name> - JunctionPath <vol_junction_path> -  
TemplateName C:\Program Files \SnapDrive\Templates\HyperVVHDxProvTemplate.xml -Size <vol_size> -  
StorageSystem <storage_system_name>
```

Similarly, you can use the `Remove-SDVolume` cmdlet to remove a volume from the SVM. When you remove a volume by using this cmdlet, the volume is dismounted, brought offline, and deleted.

Note: You cannot delete a volume in a SnapMirror relationship.

CIFS Shares

As discussed in the “Provisioning Templates” section, you can provision SMB shares by using templates for specific workloads, such as SQL Server over SMB, Hyper-V over SMB, and home directories.

The following example shows how users can provision a share for a home directory.

```
New-SdSMBshare -templatename "sd_homedir_prov_template" -Path "/homedir" -CIFSServer "fileserver"
```

Best Practices

- After provisioning a CIFS share from SnapDrive, NetApp recommends validating your environment by running the PowerShell cmdlet `debug-sdhost`. This validation enables adherence to the best practices for the workload being hosted.
- Add a DNS entry for data LIFs mapped to each CIFS server name. If multiple data LIFs are mapped to a CIFS server, add any of the data LIFs.

13 Backup and Restore Architecture of a CIFS Share Using SnapDrive for Windows

13.1 Backup and Restore Architecture

SnapManager for Hyper-V and SnapManager for SQL Server support SMB workflows. When a backup or restore operation is initiated from one of these SnapManager products, calls are passed to the SnapDrive web service proxy. Then, using the MSRPC protocol, the SnapDrive web service proxy initiates communication between the local file-share copy provider from Microsoft and the remote VSS provider.

13.2 Backup and Restore

SnapDrive supports the backup and restore of SMB shares. Multiple shares hosted across different CIFS servers can be backed up at once. Snapshot copies can be created using the `New-SDsnapshot` cmdlet. You can use this cmdlet to create Snapshot copies of a LUN or an SMB share. The following example backs up the SMB shares by creating Snapshot copies of the corresponding volumes using the Snapshot name `sql_snap`.

```
New-SdSnapshot -path "\\fileserver\sqlshare","\\fileserver\sqlshare2" -snapshot "sql_snap"
```

The `restore-SDSnapshot` cmdlet is used for restoring both LUNs and SMB shares locally or from the secondary storage (SnapVault). In the case of SMB shares, individual files and directories can also be restored. SnapDrive also supports restoring SMB shares from secondary storage such as SnapVault.

Examples of Typical Restore Operations

Example 1: Restore a Snapshot Copy on an SMB Share From a Snapshot Copy.

```
Restore-SdSnapshot -Path "\\172.17.12.101\share\files.txt" -Snapshot "snapshot_1"
```

This example restores the file named `file.txt` on SMB share `\\172.17.12.101\share` from the specified Snapshot copy `snapshot_1`.

Example 2: Restore a File Under a Subfolder of an SMB Share From a Snapshot Copy.

```
Restore-SdSnapshot -Path "\\172.17.12.101\share\dir1\file1.txt" -Snapshot snapshot_1
```

This example restores the file at `\\172.17.12.101\share\dir1\file1.txt` from the Snapshot copy `snapshot_1`.

Example 3: Restore a Directory with its Contents Under an SMB Share From a Snapshot Copy.

```
Restore-SdSnapshot -Path "\\172.17.12.101\share\folder1*" -Snapshot "snapshot_1"
```

This example restores the directory named `folder1` and its contents from the specified Snapshot copy `snapshot_1`.

Example 4: Restore Multiple Files and Directories Under an SMB Share from a Snapshot Copy.

```
Restore-SdSnapshot -Path
"\\172.17.12.101\share\file0.txt", "\\172.17.12.101\share\dir1\file1.txt", "\\172.17.12.101\share\dir2\
-Snapshot snapshot_1
```

This example restores a file named `file0.txt` under the root of the SMB share, a file named `file1.txt` under directory `dir1`, and directories named `dir2` and `dir3` and their contents from the Snapshot copy named `snapshot_1`.

Example 5: Restore a File on an SMB Share from a Snapshot Copy on the SnapVault Secondary.

```
Restore-SdSnapshot -Path "\\172.17.12.101\share\dir1\file1.txt" -Snapshot "snapshot_1" -
StorageSystem 172.17.165.29 -VolumeName vaultdest_vol
```

This example restores a file named `file1.txt` under directory `dir1` from the Snapshot copy named `snapshot_1` on the SnapVault secondary storage system.

Example 6: Restore a Snapshot Copy on a Disk from a Snapshot Copy.

```
Restore-SdSnapshot -Path E: -Snapshot "snapshot_1"
```

This example restores the Snapshot copy named `snapshot_1`.

Note: The `restore-sdsnapshot` cmdlet cannot perform file-level restores in SAN environments.

Best Practices

- In case of a backup failure, use the `Get-SdSMBShadowCopyEmsMessage` cmdlet to retrieve EMS logs from the SVM specific to the backup event. This strategy helps determine the root cause of the failure.
- If a user must turn a volume offline for any reason, verify that the volume is remounted using the correct junction path. Otherwise, backup operations are likely to fail.

13.3 Mounting and Dismounting a Snapshot Copy

Use the `Mount-SdSnapshot` cmdlet to mount a LUN or a share from the specified Snapshot copy and show them as a different set of shares having a new global unique identifier (GUID).

This cmdlet can be used in an SQL-over-SMB environment to perform backup verification. The user can mount the database and log shares from the Snapshot copy and carry out database verification. The access control lists (ACLs) on the mounted share are the same as those on the original share.

You can also use the `mount-sdsnapshot` cmdlet to mount the shares from a secondary Snapshot copy (SnapVault). To do so, you must specify the storage system and volume. Verify that the aggregate of the volume that is the source of the FlexClone operation is assigned to the SVM aggregates list. The following example mounts the share `DBShare` from the specified secondary Snapshot copy `weekly_snap`.

```
Mount-SdSnapshot -Path "\\SQLFileserver\DBShare", "\\SQLFileserver\LogShare" -snapshot
"weekly_snap" -storagesystem mirror_vserver -volume dbmirrorvolume
```

When mounting a Snapshot copy, verify that the aggregate of the volume that is the source of the mount operation is added to the SVM-assigned aggregates list option.

A share is deleted upon dismounting, and the underlying FlexClone volume is also deleted.

Note: The FlexClone volume that is created during the mount operation is not space guaranteed.

14 Creating VMs in CIFS Shares Created by SnapDrive for Windows

After using SDW to provision the SMB shares, it is a straightforward process to configure Hyper-V in Windows Server 2012 and Windows Server 2012 R2 to use these shares as storage for VM virtual disks. To do so, set the location of VHDX files (Figure 5) and VM configuration files (Figure 6).

Figure 5) VHDX location settings.

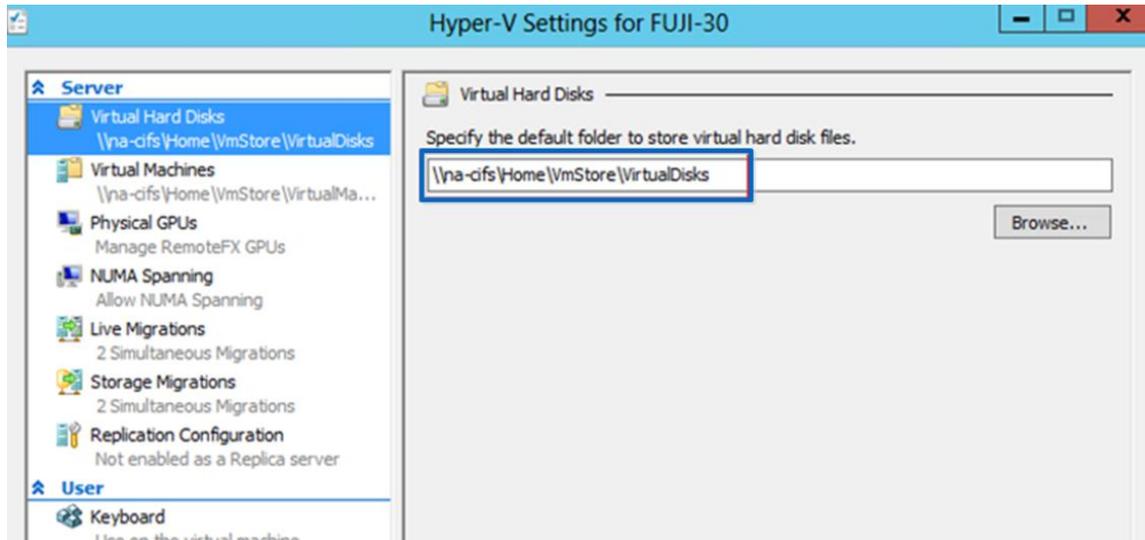
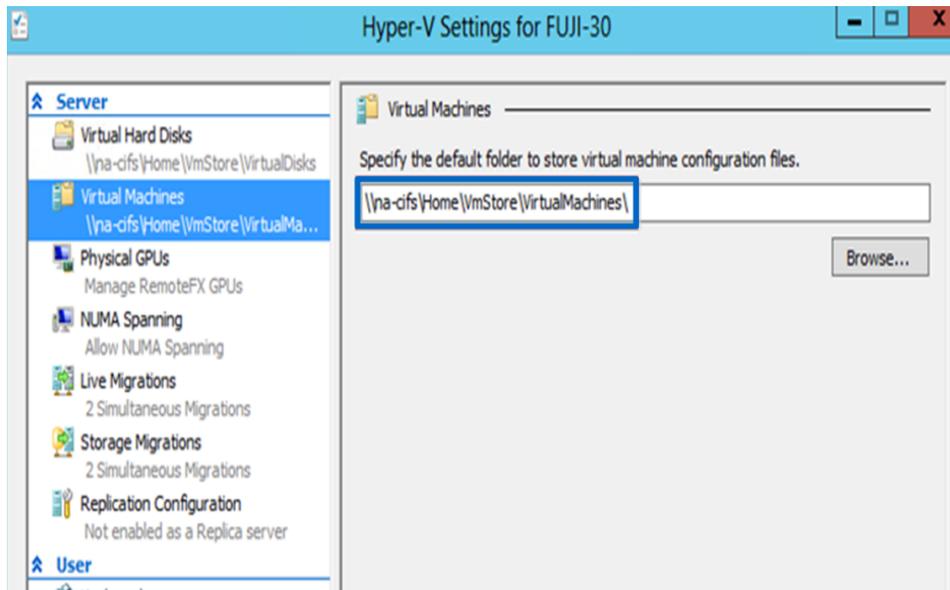


Figure 6) VM location settings.



These VMs can then be protected with SnapManager for Hyper-V by using specific backup and replication policies, and the `Get-SdVM` PowerShell cmdlet can be used to retrieve VM information.

15 SnapMirror and SnapVault

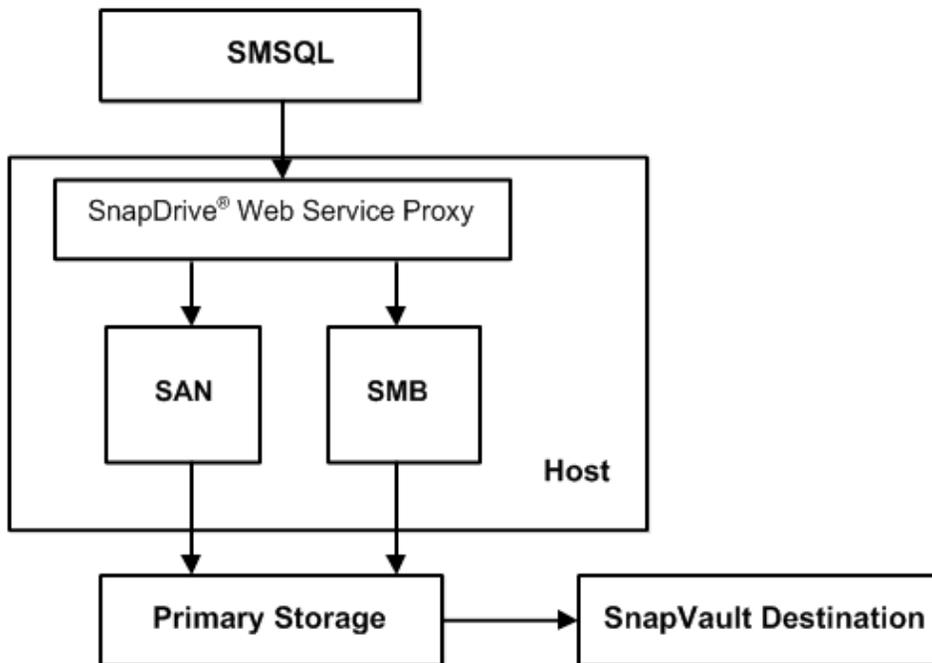
15.1 SnapMirror

SDW supports SnapMirror for volumes containing SMB shares. If the SnapMirror source volume is replicated to multiple destination volumes (a fan-out scenario), all of the destinations are updated. For best practices on SnapMirror configurations, refer to [TR-4015: SnapMirror Configuration and Best Practices Guide for Clustered Data ONTAP](#).

15.2 SnapVault

SDW introduces native SnapVault integration. Therefore, OnCommand® Protection Manager is no longer required to configure and update SnapVault datasets in environments running clustered Data ONTAP 8.2 and later. Rather, SnapVault can be configured by using PowerShell cmdlets and initiated with the SnapManager for SQL Server or SnapManager for Hyper-V GUI. Figure 7 depicts the workflow for SnapVault.

Figure 7) SnapVault workflow.



Note: In the case of Data ONTAP systems operating in 7-Mode, SnapDrive continues to require OnCommand Protection Manager for SnapVault configuration in SAN environments.

Note: Cascaded SnapVault is not supported.

To configure specific retention-period policies for the Snapshot copies, use the `Set-SdSnapMirrorPolicyRule` PowerShell cmdlet. Each volume that has a SnapVault relationship that hosts SMB shares can have a SnapVault policy associated with it. These rules help users to configure the retention periods and the actions taken after the threshold is reached. Users have the option to create custom labels depending on specific business needs.

The `Set-SdSnapMirrorPolicyRule` cmdlet can be used to set the SnapVault retention policies as well as thresholds. For example, the following command sets the SnapMirror policy rule `myWeekly` for the policy of the specified relationship.

```
Set-SdSnapMirrorPolicyRule -SourceStorageSystem vs01 -SourceStorageSystemVolume src_vol01 -
DestinationStorageSystem vs02 -DestinationStorageSystemVolume dest_vol01 -SnapMirrorLabel
myWeekly -Retention 8 -Preserve -WarnThreshold 3 -verbose -Confirm:$false
```

Before setting the policy rule, verify that the policy is created at the cluster SVM.

Users also can use the PowerShell cmdlet `set-SDsnapshot` to attach labels to Snapshot copies and then select the secondary retention bucket by specifying the appropriate label.

The following example adds the suffix label `monthly` to the specified Snapshot copy `salesdb_backup`.

```
Set-SdSnapshot -storagesystem prodvserver -volume voldb,vollog -snapshot salesdb_backup | set-
Sdsnapshot -label monthly
```

The `GetMirrorInfo` cmdlet can be used to retrieve information on the SnapVault and SnapMirror relationships that were established from a given storage system.

```
PS C:\Users\administrator.TEST>(Get-SdStorage -StorageSystem 172.17.162.61 -
GetMirrorInfo -Verbose).StorageSystemResource.volume
```

When restoring from a SnapVault storage system, verify that a valid FlexClone and CIFS license was installed on the SnapVault system. In addition, a valid CIFS server must be present in the SnapVault system.

Place the secondary CIFS server in the same domain as the primary CIFS server; if they are in different domains, there must be trust between the two domains.

Add the SnapVault and SnapMirror destination SVM credentials in the transport protocol settings. If the SnapVault destination is in the same cluster, then use ODX.

16 Restoring Snapshot Copies After Storage Live Migration

The storage live migration feature in Windows Server 2012 and Windows Server 2012 R2 enables migration of VM-related files to a different storage location without causing the VM to undergo downtime. This process is faster if the storage system supports ODX. As discussed in the section “Features in Clustered Data ONTAP,” Data ONTAP 8.2 and later support ODX.

When you initiate storage migration of a VM from one SMB share to another SMB share or LUN within the same volume or on a different volume, Windows Server 2012 queries the storage system on whether it is copy-offload enabled. If it is, Windows Server then offloads the copy activity to the storage system.

Note: It is best to avoid SnapDrive operations during storage live migration because they might corrupt the VM.

Note: After performing storage migration of a VM from one share to another share hosted on a different volume, the SnapMirror and SnapVault relationships must be reestablished with the new destination. In addition, previously existing Snapshot copies cannot be restored from the SnapVault storage system.

Version History

Version	Date	Document Version History
Version 1.0	August 2013	Initial release

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

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-4218-0115

