Technical Report

# Microsoft Hyper-V over SMB 3.0 with Clustered Data ONTAP: Best Practices

Pavel Lobanov, John Reed, NetApp
May 2013 | TR-4172

## Abstract

This technical report provides guidelines and best practices for integrated architecture and implementation of Microsoft® Hyper-V™ with NetApp® storage solutions utilizing the SMB 3.0 NAS protocol. The NetApp technologies discussed in this technical report provide customers with an integrated storage solution that is cost effective, operationally efficient, flexible, and environmentally friendly.

TABLE OF CONTENTS


1 Purpose and Scope ..... 4

2 Clustered Data ONTAP 8.2 ..... 4

2.1 SMB 3.0 Protocol in Clustered Data ONTAP 8.2 ..... 4

2.2 Clustered Data ONTAP 8.2 Setup ..... 6

2.3 SMB 3.0 Protocol Negotiation Feature ..... 6

2.4 General Considerations of Clustered Data ONTAP for SMB 3.0 ..... 8

3 CIFS Vserver and Share Setup Configuration ..... 8

3.1 Root and Data Volume Settings ..... 8

3.2 Data and Management LIF Settings ..... 8

3.3 SMB 3.0 Settings ..... 8

3.4 ODX Settings ..... 9

3.5 Remote VSS Settings (Shadow Copy Feature VSS) ..... 10

3.6 Automatic Node Referral Settings ..... 10

3.7 Creating Continuously Available File Shares ..... 11

3.8 Creating Network Interface Failover Groups ..... 11

4 Microsoft Windows Server 2012 Improvements ..... 12

4.1 SMB 3.0 ..... 12

4.1 Hyper-V Improvements ..... 12

5 Microsoft Hyper-V Windows Server 2012 Settings ..... 12

5.1 Hyper-V Server Farms ..... 13

5.2 Types of Virtual Disks ..... 14

5.3 Fast Provisioning of Virtual Disks Using the ODX Feature ..... 14

6 SnapManager for Hyper-V 2.0 (SMHV 2.0) ..... 15

6.1 Remote VSS ..... 15

6.2 SMHV 2.0 Components ..... 15

6.3 SMHV Windows PowerShell CmdLets ..... 16

7 SMHV 2.0 Configuration and Operations ..... 16

7.1 Connecting SMHV to NetApp Storage System ..... 16

7.2 Creating Datasets and Backup Policies ..... 17

7.3 Virtual Machine Backups and Restores ..... 17

**References** ..... **19**

**Version History** ..... **19**

## LIST OF FIGURES

Figure 1) SMB negotiate request. ..... 6
Figure 2) SMB negotiate response. ..... 7
Figure 3) Tree connect response. ..... 7
Figure 4) VHDX location settings. ..... 12
Figure 5) VM location settings. ..... 13
Figure 6) Hyper-V VM move options. ..... 14
Figure 7) NetApp remote VSS components. ..... 15
Figure 8) Storage connection settings. ..... 17
Figure 9) Manual backup of VMs. ..... 17
Figure 10) Restoring VMs. ..... 18
Figure 11) Restore options. ..... 18

# 1 Purpose and Scope

NetApp has been at the forefront of solving complex business problems with its innovative technology breakthroughs and end-to-end solutions approach. One of NetApp's strongest areas in solving these problems is the use of virtualization throughout the data center. Server virtualization is a major component in this space, and Microsoft is a lead player with its server virtualization solutions. This technical report provides detailed guidance on how to architect and implement Microsoft server virtualization solutions on NetApp storage using the SMB 3.0 NAS protocol with clustered Data ONTAP® 8.2. It describes the use of and best practices for setting up a Windows® 2012 Hyper-V environment with SnapManager® for Hyper-V (SMHV) 2.0 to protect VMs running on SMB 3.0 file share hosting on NetApp clustered Data ONTAP.

This technical report is not intended to be a definitive implementation or solutions guide. Expertise might be required to solve specific deployments. Contact your local NetApp sales representative to speak with one of our Microsoft Hyper-V solutions experts. NetApp is dedicated to helping you transform your data center to help your business go further, faster.

# 2 Clustered Data ONTAP 8.2

Clustered Data ONTAP 8.2 provides a complete solution for NetApp customers to deploy a virtualized environment and protect virtual machines (VMs) running on file-level data storage provided by the SMB 3.0 protocol.

## 2.1 SMB 3.0 Protocol in Clustered Data ONTAP 8.2

One of the major components added to clustered Data ONTAP 8.2 is support for the SMB 3.0 NAS protocol, which enables NetApp customers to utilize the SMB 3.0 features introduced with Windows Server® 2012. With these new features, clustered Data ONTAP can be used 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 to support continuously available file shares and Hyper-V storage are:

- Persistent handles (continuously available file shares)
- Witness protocol
- Clustered client failover
- Scale-out awareness
- Offloaded data transfer (ODX)
- Remote VSS

### Persistent Handles (Continuously Available File Shares)

To enable continuous availability on a file share, the SMB client opens a file on behalf of the application, such as a VM running on a Hyper-V host, and requests persistent handles for the 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 case of a planned move of file share resources from one node to another, or in case of failure of the node, the SMB client will reconnect to an active and available node and will reopen the file using persistent handles. The application/VM running on the SMB client computer does not experience any failures or errors during this operation. From a VM perspective, it appears the I/O operations to virtual disk were delayed for a small amount of time, similar to brief loss of connectivity to the disk, but no disruption is noticed.

### Witness Protocol

When an SMB server node fails, the SMB client usually relies on the 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

since the virtual disk is a critical resource, the VM running on a Hyper-V server needs faster detection of network resources failing over. Witness protocol significantly improves the SMB client reconnect time.

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

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 a reconnect to the running node immediately, minimizing the time the client is disconnected from the resource. For VMs with virtual disks stored on such SMB shares, disk disconnection time is reduced to the point where the VM would 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 LIF on each node in the cluster in every Vserver. Note also that witness protocol only comes into play for continuously available shares.

## Clustered Client Failover (CCF)

To increase redundancy in a VM environment, Hyper-V servers should be placed into a Microsoft failover cluster. When the Hyper-V server node running a VM fails, the VM is live-migrated/moved to another node. Before CCF with SMB 3.0, a VM moving to another cluster node was considered as a new application instance. New application instances connecting to files already open on file shares have to wait until the TCP timeout is over and the file handle is closed. CCF gives the VM the ability to open a virtual disk file on a file share and provide 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, and the SMB client can then reconnect to the previously open file.

## Scale-Out Awareness

Clustered Data ONTAP is scale-out by design and provides the ability to serve data from multiple nodes. It brings additional data redundancy on 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 get to the same data.

## Offloaded Data Transfer (ODX, Copy Offload)

Although this feature is not required to run a Hyper-V workload over SMB 3.0, in typical deployments where the customer needs to provision multiple VMs, this feature can drastically improve VM deployment time. The main advantage of this feature is that it is transparent to client machines, and no actual data is put over the network during file copy operations. Clustered Data ONTAP provides different mechanisms on the back end to copy data blocks. In the case of a single volume serving a file share, NetApp uses its SIS-clone functionality, which eliminates the data copy process by creating only pointers. This will speed up back-end operations and improve copy performance with ODX on the NetApp platform when compared to ODX implementations in other storage arrays. When data is copied within the cluster, but outside the volume, the process remains offloaded, and no traffic will go through the client or the network.

## Remote VSS

Volume Shadow Copy Service (VSS) is a framework that provides coordination of application I/O and physical storage on the same server and allows creation of application-consistent Snapshot™ copies of the storage. Microsoft Windows Server 2012 extends the functionality of VSS to multiple servers. For instance, 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 backup Snapshot copies of the storage for applications running remotely on the storage server. Clustered Data ONTAP 8.2 extends 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 a Windows Server 2012 machine. From a VSS perspective, the NetApp array acts in the same way as a Windows File Server.

## 2.2 Clustered Data ONTAP 8.2 Setup

There are no special requirements for a clustered Data ONTAP 8.2 setup to utilize SMB 3.0 features. By default, clustered Data ONTAP 8.2 supports all versions of SMB, including 1.0, 2.0, 2.1, and 3.0. Most of the usual applications of SMB protocol, such as user file sharing, can work on pre-3.0 SMB protocols and might not benefit from the additional features of the SMB 3.0 protocol. The Hyper-V workload requires the SMB 3.0 protocol and some of its additional features, such as continuously available shares. Taking into consideration the additional overhead of storing and replicating persistent handle information between nodes in an HA pair to support features such as continuously available shares, it is strongly recommended to utilize only continuously available shares for Hyper-V over SMB workloads.

## 2.3 SMB 3.0 Protocol Negotiation Feature

SMB feature negotiation between client and server starts when a client machine initiates a connection to a CIFS/SMB file share. During the SMB negotiation request (see Figure 1), the client tells the server which dialects (versions) of SMB it supports. This section shows the negotiation process (using packet trace pictures) and how dialect and features are confirmed. If troubleshooting is required, customers can use packet traces to confirm that the correct dialect and features are being used.

**Figure 1) SMB negotiate request.**

Trace Session 1 : Analysis Grid

| MessageNum | Timest | Time | Source | Destination | Module | Summary |
|---|---|---|---|---|---|---|
| 26 | 03/14… | 0 | 97.1.1.116 | 97.1.1.21 | TCP | Flags: CE....S., Port: 49998 - 4 |
| 27 | 03/14… | 0 | 97.1.1.21 | 97.1.1.116 | TCP | Flags: ...A..S., Port: 445 - 499 |
| 28 | 03/14… | 0 | 97.1.1.116 | 97.1.1.21 | TCP | Flags: ...A...., Port: 49998 - 4 |
| 29 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Negotiate, ClientGuid = f3e3b7 |
| 29 | 03… | 0. | 97.1.1.116 | 97.1.1.21 | SMB2 | C Negotiate |
| 33 | 03… | 0 | 97.1.1.21 | 97.1.1.116 | SMB2 | R Negotiate, Revision = SMB3, |
| 34 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Session Setup, SessionFlags = |
| 40 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Tree Connect, Path = \\na-cifs |
| 45 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O IOCTL Failure, FID = 0xFFFFFFF |
| 50 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Create, Name = na-cifs\Home\Vm |
| 54 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Close, FID = 0x0, Status = Suc |

Details

| Name | Value | Type | Bit Offset | Bit Length |
|---|---|---|---|---|
| Header | SMB2PacketHea… | Type | 0 | 512 |
| Request | SMB2Negotiate… | Type | 512 | 336 |
| StructureSize | 36 | UInt16 | 512 | 16 |
| DialectCount | 3 | UInt16 | 528 | 16 |
| SecurityMode | SMB2Negotiate… | Type | 544 | 16 |
| Reserved | 0 | UInt16 | 560 | 16 |
| Capabilities | SMB2GlobalCap… | Enum | 576 | 32 |
| ClientGuid | f3e3b795-8cde… | Guid | 608 | 128 |
| ClientStartTime | 1/1/1601 0:0:… | Type | 736 | 64 |
| Dialects | [514,528,768] | Array | 800 | 48 |
| [0] | 514 | UInt16 | 202 | |
| [1] | 528 | UInt16 | 210 | |
| [2] | 768 | UInt16 | 300 | |

The server responds with the highest common version of the SMB protocol (see Figure 2).

**Figure 2) SMB negotiate response.**

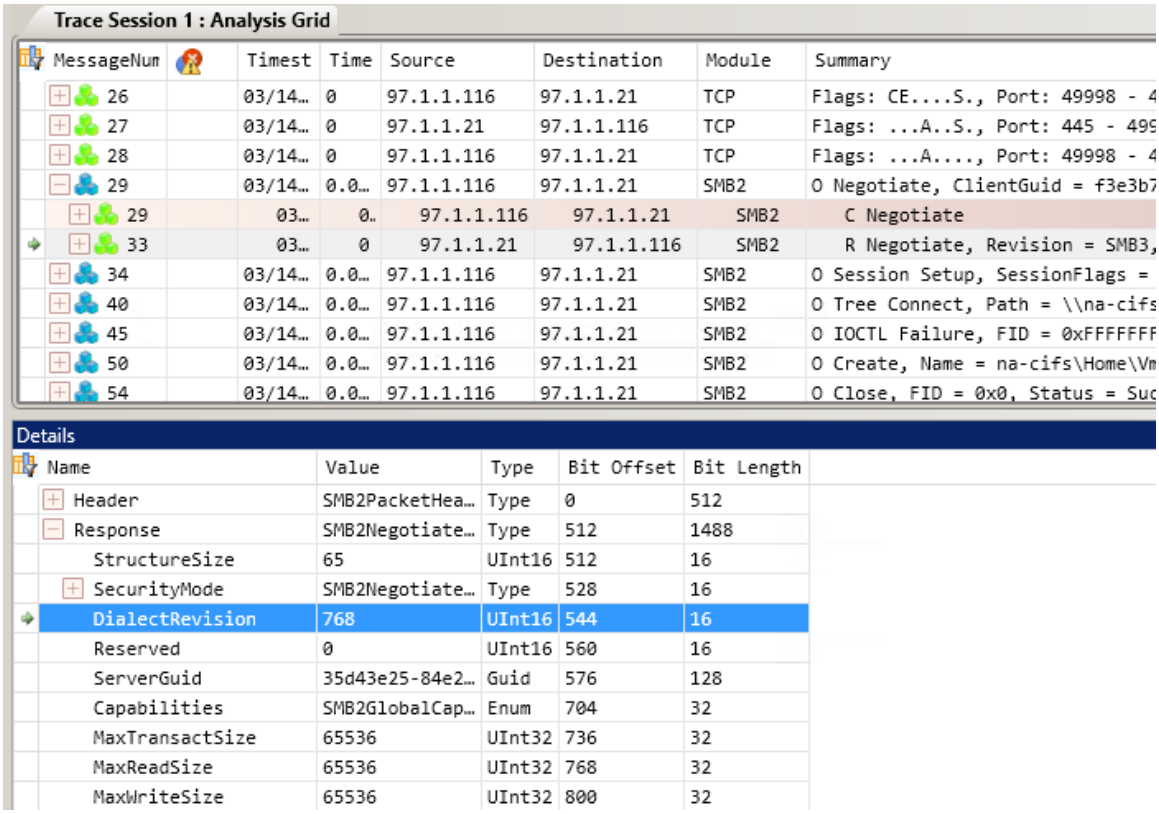

Trace Session 1 : Analysis Grid

| MessageNum | Timest | Time | Source | Destination | Module | Summary |
|---|---|---|---|---|---|---|
| 26 | 03/14… | 0 | 97.1.1.116 | 97.1.1.21 | TCP | Flags: CE....S., Port: 49998 - 4 |
| 27 | 03/14… | 0 | 97.1.1.21 | 97.1.1.116 | TCP | Flags: ...A..S., Port: 445 - 499 |
| 28 | 03/14… | 0 | 97.1.1.116 | 97.1.1.21 | TCP | Flags: ...A...., Port: 49998 - 4 |
| 29 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Negotiate, ClientGuid = f3e3b7 |
| 29 | 03… | 0. | 97.1.1.116 | 97.1.1.21 | SMB2 | C Negotiate |
| 33 | 03… | 0 | 97.1.1.21 | 97.1.1.116 | SMB2 | R Negotiate, Revision = SMB3, |
| 34 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Session Setup, SessionFlags = |
| 40 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Tree Connect, Path = \\na-cifs |
| 45 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O IOCTL Failure, FID = 0xFFFFFFF |
| 50 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Create, Name = na-cifs\Home\Vm |
| 54 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Close, FID = 0x0, Status = Suc |

Details

| Name | Value | Type | Bit Offset | Bit Length |
|---|---|---|---|---|
| Header | SMB2PacketHea… | Type | 0 | 512 |
| Response | SMB2Negotiate… | Type | 512 | 1488 |
| StructureSize | 65 | UInt16 | 512 | 16 |
| SecurityMode | SMB2Negotiate… | Type | 528 | 16 |
| DialectRevision | 768 | UInt16 | 544 | 16 |
| Reserved | 0 | UInt16 | 560 | 16 |
| ServerGuid | 35d43e25-84e2… | Guid | 576 | 128 |
| Capabilities | SMB2GlobalCap… | Enum | 704 | 32 |
| MaxTransactSize | 65536 | UInt32 | 736 | 32 |
| MaxReadSize | 65536 | UInt32 | 768 | 32 |
| MaxWriteSize | 65536 | UInt32 | 800 | 32 |

In response to the tree connect request from the client, the NetApp SMB 3.0 CIFS server responds with a set of supported features, as seen in Figure 3.

**Figure 3) Tree connect response.**

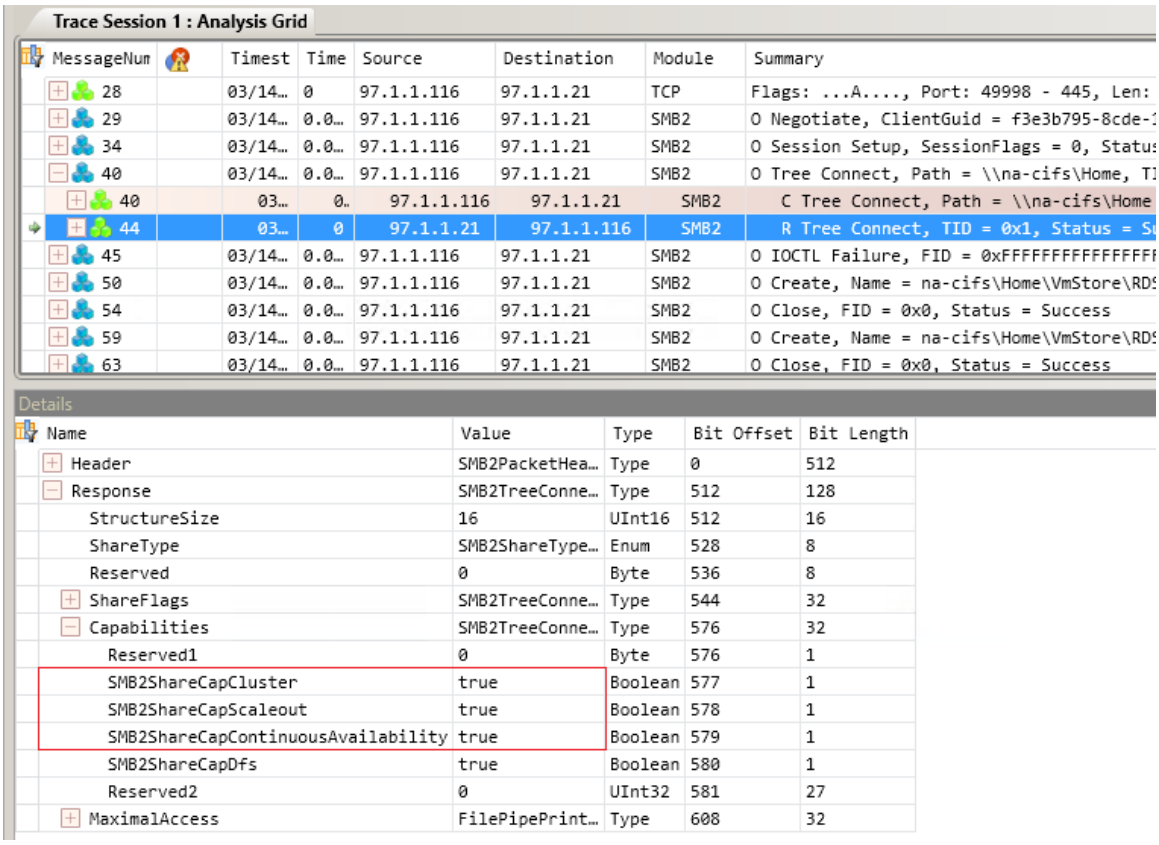

Trace Session 1 : Analysis Grid

| MessageNum | Timest | Time | Source | Destination | Module | Summary |
|---|---|---|---|---|---|---|
| 28 | 03/14… | 0 | 97.1.1.116 | 97.1.1.21 | TCP | Flags: ...A...., Port: 49998 - 445, Len: |
| 29 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Negotiate, ClientGuid = f3e3b795-8cde-1 |
| 34 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Session Setup, SessionFlags = 0, Status |
| 40 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Tree Connect, Path = \\na-cifs\Home, T |
| 40 | 03… | 0. | 97.1.1.116 | 97.1.1.21 | SMB2 | C Tree Connect, Path = \\na-cifs\Home |
| 44 | 03… | 0 | 97.1.1.21 | 97.1.1.116 | SMB2 | R Tree Connect, TID = 0x1, Status = Su |
| 45 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O IOCTL Failure, FID = 0xFFFFFFFFFFFFFFFF |
| 50 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Create, Name = na-cifs\Home\VmStore\RD5 |
| 54 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Close, FID = 0x0, Status = Success |
| 59 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Create, Name = na-cifs\Home\VmStore\RD5 |
| 63 | 03/14… | 0.0… | 97.1.1.116 | 97.1.1.21 | SMB2 | O Close, FID = 0x0, Status = Success |

Details

| Name | Value | Type | Bit Offset | Bit Length |
|---|---|---|---|---|
| Header | SMB2PacketHea… | Type | 0 | 512 |
| Response | SMB2TreeConne… | Type | 512 | 128 |
| StructureSize | 16 | UInt16 | 512 | 16 |
| ShareType | SMB2ShareType… | Enum | 528 | 8 |
| Reserved | 0 | Byte | 536 | 8 |
| ShareFlags | SMB2TreeConne… | Type | 544 | 32 |
| Capabilities | SMB2TreeConne… | Type | 576 | 32 |
| Reserved1 | 0 | Byte | 576 | 1 |
| SMB2ShareCapCluster | true | Boolean | 577 | 1 |
| SMB2ShareCapScaleout | true | Boolean | 578 | 1 |
| SMB2ShareCapContinuousAvailability | true | Boolean | 579 | 1 |
| SMB2ShareCapDfs | true | Boolean | 580 | 1 |
| Reserved2 | 0 | UInt32 | 581 | 27 |
| MaximalAccess | FilePipePrint… | Type | 608 | 32 |

Figure 3 shows that the NetApp CIFS server supports these three features:

- Cluster
- Scale-out
- Continuous availability

Based on the set of features supported by the CIFS server, the client uses features such as persistent handles and witness protocol.

### 2.4 General Considerations of Clustered Data ONTAP for SMB 3.0

When setting up clustered Data ONTAP 8.2 for using continuously available file shares to host VHDX disk images of VMs, consider the following:

- Persistent handles work only between nodes in an HA pair.
- Witness protocol works only between nodes in an HA pair.
- Continuously available file shares are only supported for Hyper-V workloads.
- ODX is supported with clustered Data ONTAP 8.2 and works across protocols. Copying data between a file share and iSCSI or an FCP-attached LUN utilizes ODX.
- Connectivity between Hyper-V hosts and the NetApp array is recommended on a 10GB network if one is available. In case of 1GB network connectivity, NetApp recommends creating an interface group consisting of multiple 1GB ports.
- CIFS and FlexClone® (required by remote VSS of SMHV) licenses should be installed.
- Time settings on nodes in the cluster should be set up accordingly. Network Time Protocol (NTP) should be used if the NetApp CIFS server has to participate in the Windows Active Directory® (AD) domain.

## 3 CIFS Vserver and Share Setup Configuration

There are several considerations when setting up an SMB 3.0 CIFS Vserver for use as storage with Windows Server 2012 Hyper-V. Some of these limitations come from the Windows AD domain joined CIFS file server, and some are defined by SnapManager for Hyper-V. The minimum Microsoft OS versions supporting SMB 3.0 are Windows Server 2012 and Windows 8.

### 3.1 Root and Data Volume Settings

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

### 3.2 Data and Management LIF Settings

At least one data LIF should be created per node for every Vserver in the cluster. The data LIF should not be configured to “AutoRevert.” Each LIF’s IP address should have an entry in DNS, and no NetBIOS aliases are allowed for DNS entries. Network interface failover groups could be configured to specify network ports to which the LIF can be moved. Also, SMHV requires one additional management LIF for the Vserver.

### 3.3 SMB 3.0 Settings

Though SMB 3.0 protocol is enabled by default, it can be checked, enabled, or disabled 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 check if 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::*>
```

## 3.4 ODX Settings

To utilize 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 per-Vserver basis, by default. To check if ODX 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::*>
```

## 3.5 Remote VSS Settings (Shadow Copy Feature VSS)

This feature should be enabled if SMHV is to be deployed on Hyper-V servers to protect VMs. 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::*>
```

## 3.6 Automatic Node Referral Settings

The Microsoft Hyper-V host relies heavily on Kerberos authentication, which cannot be utilized with NetApp IP-based automatic node referral. By default, node referrals are disabled, but when deploying Hyper-V over SMB, be sure to verify this 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::*>
```

## 3.7 Creating Continuously Available File Shares

In order to provide continuous availability, file shares on the NetApp CIFS Vserver should be set with the continuously available property. Continuously available shares are only supported for Hyper-V VM repositories and should not be created for general SMB shares, including for home folder deployments. Share properties such as homedirectory, branchcache, access-based enumeration, or attribute caching should not be set.

To set properties for file share, use the command `vserver cifs share properties add -vserver <vserver name> -share-name <share name> -share-properties continuously-available`.

```
Vespus::*> vserver cifs share properties add -vserver nacifs -share-name Home -share-
properties continuously-available

Vespus::*>
```

To check if the continuously available properties of the share are set, use the command `vserver cifs share properties show`.

```
Vespus::*> vserver cifs share properties show
Vserver          Share                    Properties
---------------  -----------------------  ----------------------------
nacifs           admin$                   browsable
nacifs           fdsa                     oplocks
                                          browsable
                                          changenotify
nacifs           Home                     oplocks
                                          browsable
                                          changenotify
                                          continuously-available
nacifs           ipc$                     browsable
nacifs           odx                      oplocks
                                          browsable
                                          changenotify
Vespus::*>
```

To allow Hyper-V to operate on file share, share and NTFS permissions should allow the Hyper-V server account to have full control.

## 3.8 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_e1a -node Vespus-01
-port e1a

Vespus::> network interface failover-groups create -failover-group nacifs_e1a -node Vespus-02
-port e1a

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.
```

# 4 Microsoft Windows Server 2012 Improvements

Microsoft Windows Server 2012 is new server platform for data center and cloud-based services. It introduces a set of new and improved features related to networking, storage, and virtualization capabilities along with improved server management and management automation. This section describes several major features related to interoperability between Microsoft and NetApp technologies, including improvements in networking (SMB 3.0 protocol) and virtualization (Hyper-V).

## 4.1 SMB 3.0

The SMB 3.0 protocol was introduced with the Windows 8 platform and has the full support of Microsoft Windows Server 2012 and Windows 8 client OSs. It enables applications such as Hyper-V to store data on file shares utilizing existing network infrastructure, instead of using the block-level storage of Fibre Channel and iSCSI.

## 4.1 Hyper-V Improvements

One of the major features of Hyper-V in conjunction with SMB 3.0 is the ability to store VM virtual disks and configuration files on remote file shares instead of local drives or cluster shared volumes (CSVs). CSVs are generally created on LUNs stored over block protocols such as iSCSI and FCP. NetApp has traditionally provided a robust storage solution for Hyper-V over block protocols. With SMB 3.0 this opens up opportunities for NetApp to provide full multiprotocol storage solutions, including SMB 3.0 file shares along with SAN-attached block-level storage within the same unified architecture. Customers can rapidly deploy Hyper-V on NetApp SMB3 file shares connected over Ethernet with minimal changes required to network data center design.

# 5 Microsoft Hyper-V Windows Server 2012 Settings

Configuring Hyper-V in Windows Server 2012 to use remote CIFS/SMB file shares as storage for VM virtual disks is simple and requires setting the location of VHDX files (see Figure 4) and VM configuration files (see Figure 5).

**Figure 4) VHDX location settings.**

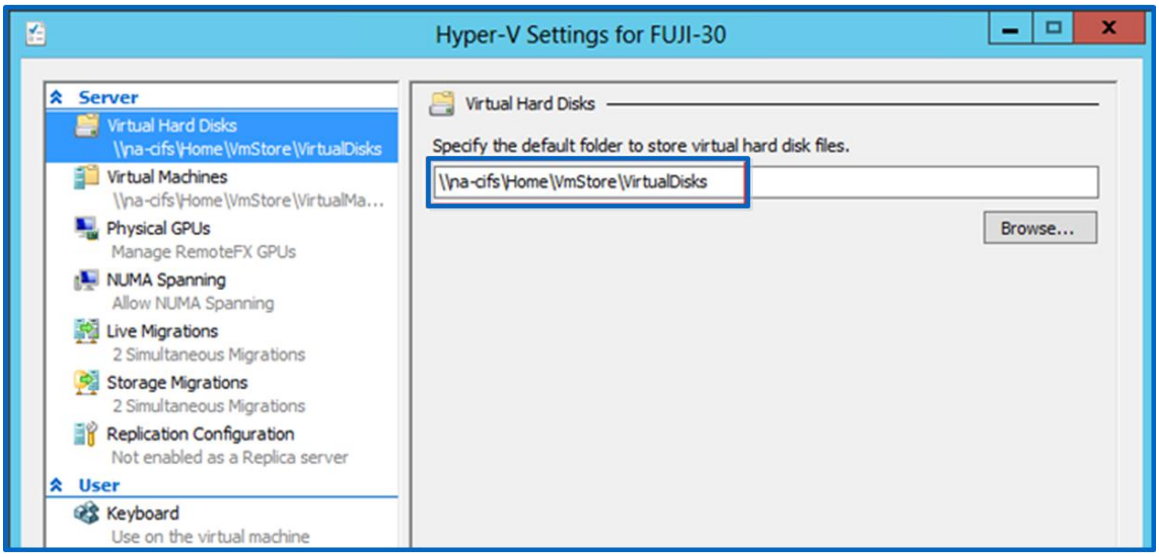

Figure 5) VM location settings.

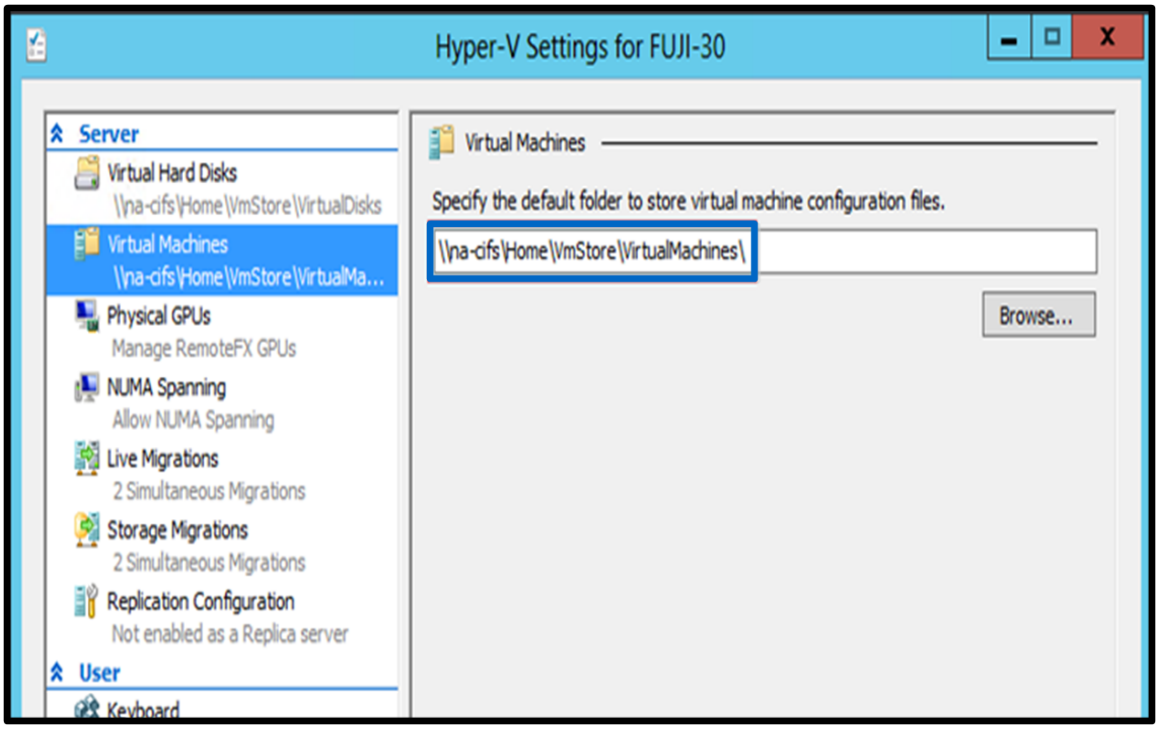


## 5.1 Hyper-V Server Farms

Previously Microsoft recommended using Microsoft failover clustering of Hyper-V servers with SAN-attached CSVs for hosting VM virtual disks. Storing virtual disks on remote file shares no longer requires CSVs for a redundant Hyper-V failover cluster. It eliminates the need of SAN infrastructure in addition to simplifying network configuration, while maintaining the same level of redundancy at the Hyper-V host level.

With Windows Server 2012, Microsoft has introduced the new feature of moving VMs between nonclustered Hyper-V hosts (see Figure 6). By using remote file shares hosting VM virtual disks, only VM information can be moved between Hyper-V hosts, leaving the VM virtual disks on the same storage accessible by multiple Hyper-V hosts. This speeds up the entire process of moving VMs and eliminates the need for a Microsoft failover cluster for less critical VMs.

Figure 6) Hyper-V VM move options.

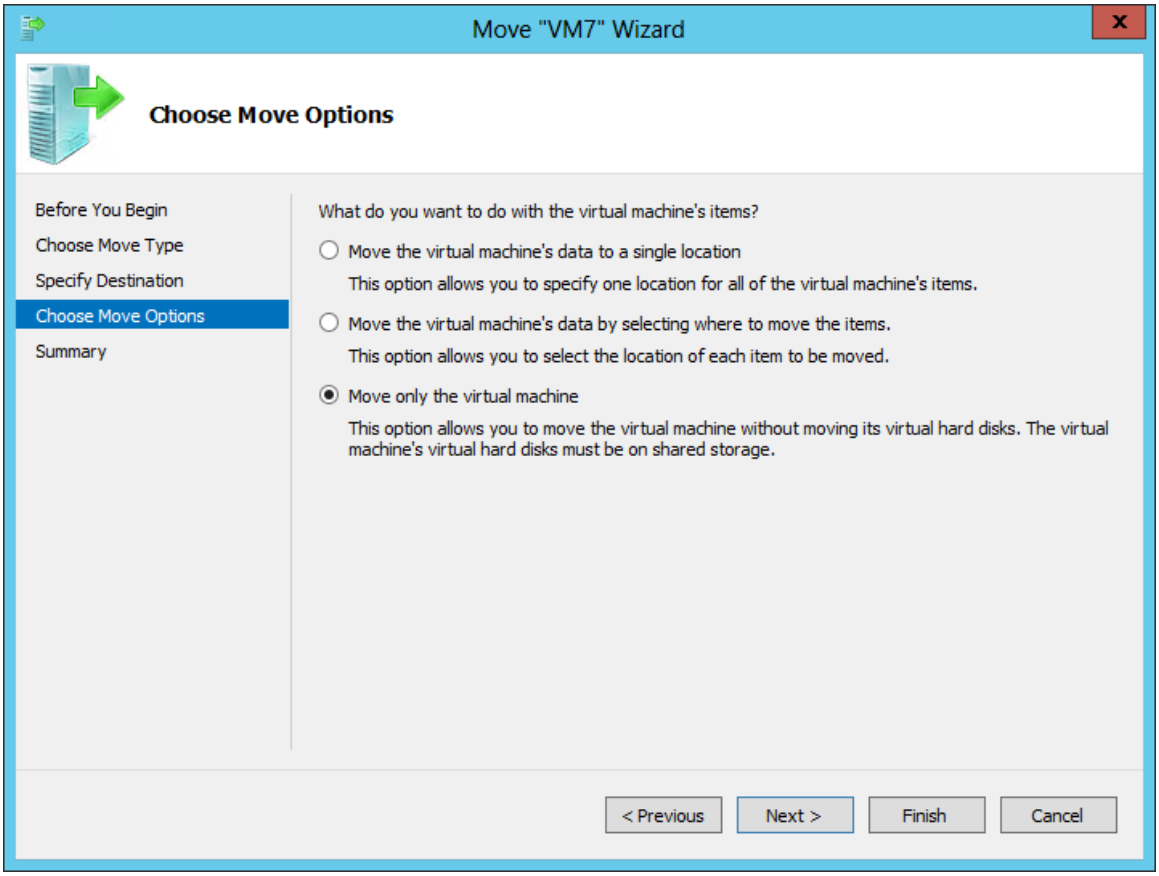


## 5.2 Types of Virtual Disks

Microsoft Hyper-V in Windows Server 2012 supports Windows Server 2008 R2 VHD styles:

- **Fixed VHD.** Historically this is the only type of VHD supported by NetApp.
- **Dynamic VHD.** Based on the structure of this VHD, it has a misalignment problem that can cause performance issues with a NetApp array.
- **Differencing VHD.** By nature, this type is less performant and used only for user desktops.

These are the new types of VHDXs implemented in Windows Server 2012:

- **Fixed VHDX.** This type is similar to the VHD format and is still recommended as the primary choice for virtual disks stored on a NetApp array. It can be thin provisioned by NetApp using FlexClone and deduplication features of clustered Data ONTAP.
- **Dynamic VHDX.** Microsoft has fixed the misalignment problems of this type of virtual disks, and performance has greatly improved and is close to the performance of a fixed VHDX.
- **Differencing VHDX.** By nature, this type is less performant and used only for user desktops.

Given the preceding information, NetApp still recommends using a fixed VHD. There are no restrictions for using any type of VHDX virtual disks.

## 5.3 Fast Provisioning of Virtual Disks Using the ODX Feature

The ODX feature in clustered Data ONTAP allows making copies of master VHDXs by simply copying a master VHDX file hosted by a NetApp array. Since an ODX-enabled copy does not put any data on the network wire, the copy process happens on the NetApp array and as a result can be up to six to eight times faster on a 10GB network. General considerations for fast provisioning include master SysPreped images stored on file shares and regular copy processes initiated by the Hyper-V host machine.

# 6 SnapManager for Hyper-V 2.0 (SMHV 2.0)

This section describes remote VSS capabilities of SMHV v2.0. Detailed information about SMHV settings, configuration, and use cases will be provided in a separate technical report.

The new version of SnapManager for Hyper-V (SMHV) is remote VSS enabled, meaning that SMHV can detect that a VM's configuration and virtual disk are stored on remote file shares hosted by a NetApp array. During a backup process, SnapManager coordinates I/O activity on local Hyper-V hosts utilizing local VSS and Snapshot captures of the NetApp volume hosting the remote file share. It does this while making sure that all virtual disks are in an application-consistent state. SMHV v2.0 is aware of clustered Hyper-V hosts and coordinates backup activities between all Hyper-V hosts in a Microsoft failover cluster. Currently SMHV v2 supports 16 Hyper-V nodes and 2,000 VMs.

## 6.1 Remote VSS

Remote VSS features allow coordination of I/O operations between the host generating the I/O and the remote host/appliance controlling the storage. The NetApp remote VSS components are described in Figure 7.

**Figure 7) NetApp remote VSS components.**

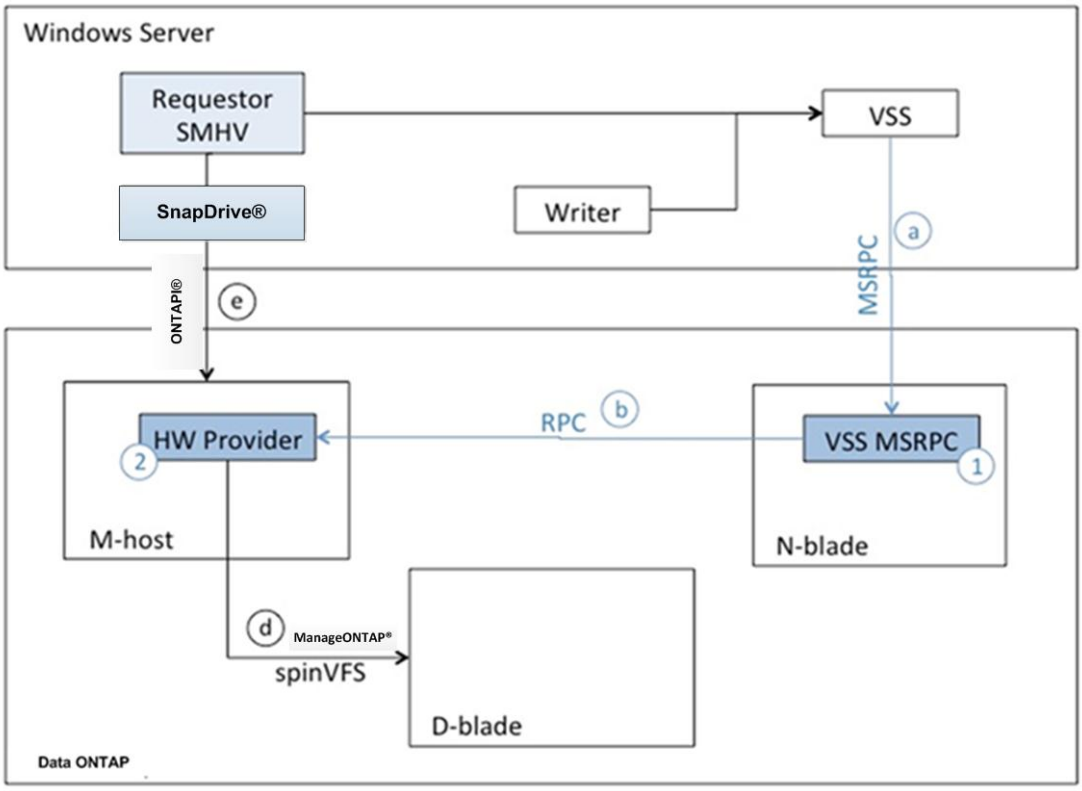


S

## 6.2 SMHV 2.0 Components

SnapManager for Hyper-V v2.0 installs the next components on the Hyper-V host:

- SMHV service
- SMHV GUI
- SnapIntegrator™ service
- SMHV Windows PowerShell™ module

## 6.3 SMHV Windows PowerShell CmdLets

Microsoft is improving Hyper-V as its virtualization platform by using the following cmdlets:

- Get-SiCifsShadowCopyEmsMessage
- Get-SiInfo
- Get-SiSnapMirror
- Get-SiSnapshot
- Get-SiStorage
- Get-SiStorageConnectionSetting
- Get-SiVM
- Invoke-SiEmsAutosupportLog
- Invoke-SiSnapMirrorUpdate
- New-SiCIFSShare
- New-SiSnapshot
- New-SiVolume
- Remove-SiSnapshot
- Remove-SiStorageConnectionSetting
- Rename-SiSnapshot
- Restore-SiSnapshot
- Set-SiStorageConnectionSetting

# 7 SMHV 2.0 Configuration and Operations

After SMHV is installed, it requires initial configuration, which includes:

- Report settings
- Notification settings
- SnapInfo settings
- Storage connection settings

All configuration settings can be done by running the configuration wizard.

## 7.1 Connecting SMHV to NetApp Storage System

The user has to provide the Vserver management LIF IP address and Vserver account with administrative privileges to connect SMHV to a NetApp storage system. Figure 8 shows the storage connection settings.

Figure 8) Storage connection settings.

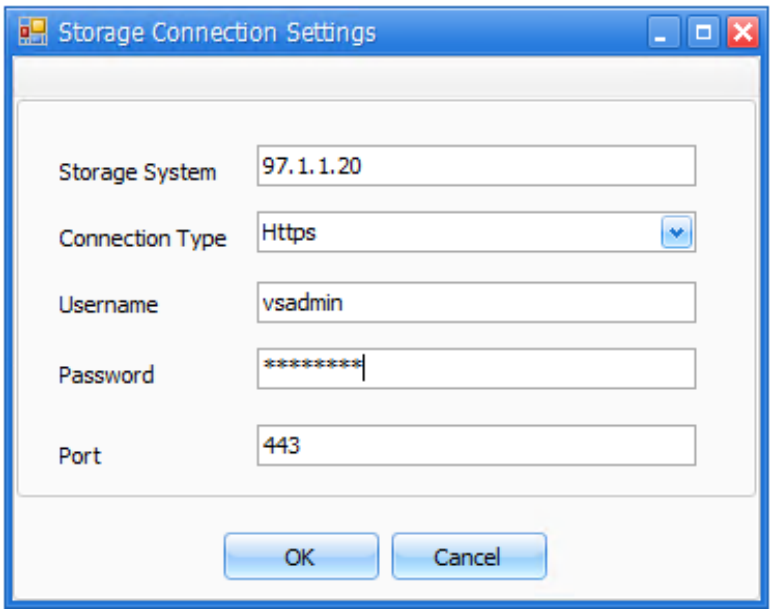


## 7.2 Creating Datasets and Backup Policies

There are two wizards that help create datasets containing a set of VMs that need to be protected and add a backup policy for each dataset. The backup policy for the dataset specifies the type of backup policy, such as application consistent or crash consistent; the backup schedule; and the retention policy. If backups are to be stored remotely, the option to update SnapMirror® after successful backup is available.

## 7.3 Virtual Machine Backups and Restores

VMs are backed up as a set by scheduled backup jobs or are manually initiated. For manual backups, choose the backup option, as seen in Figure 9.

Figure 9) Manual backup of VMs.

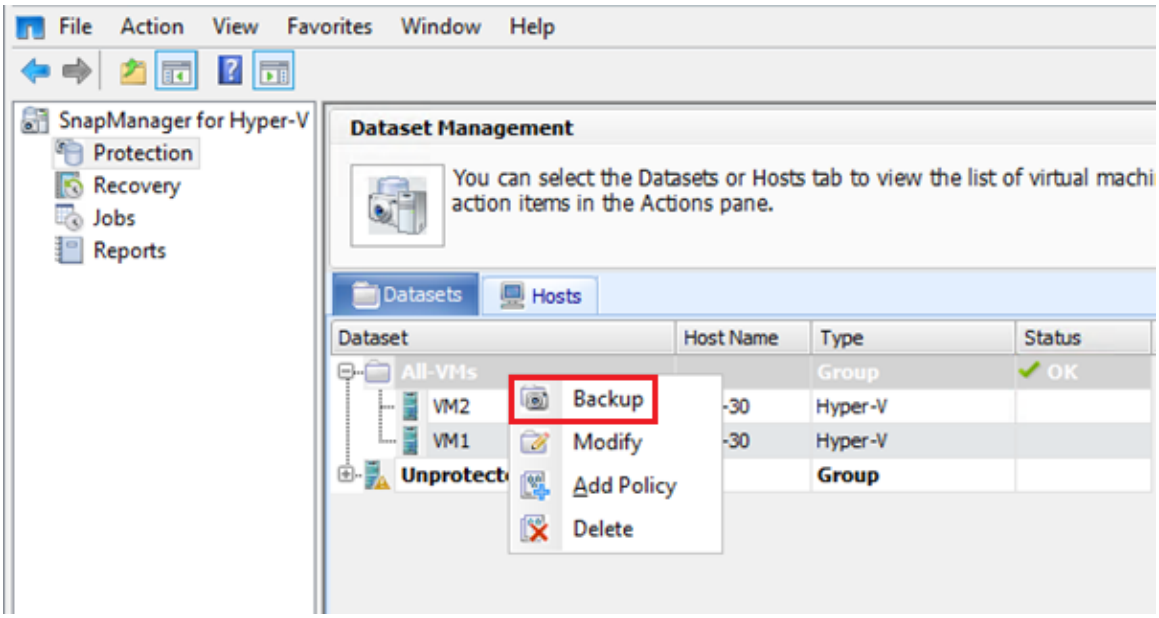


In case of a disaster, the VM can be restored from a backup Snapshot copy in a matter of seconds (see Figure 10).

Figure 10) Restoring VMs.

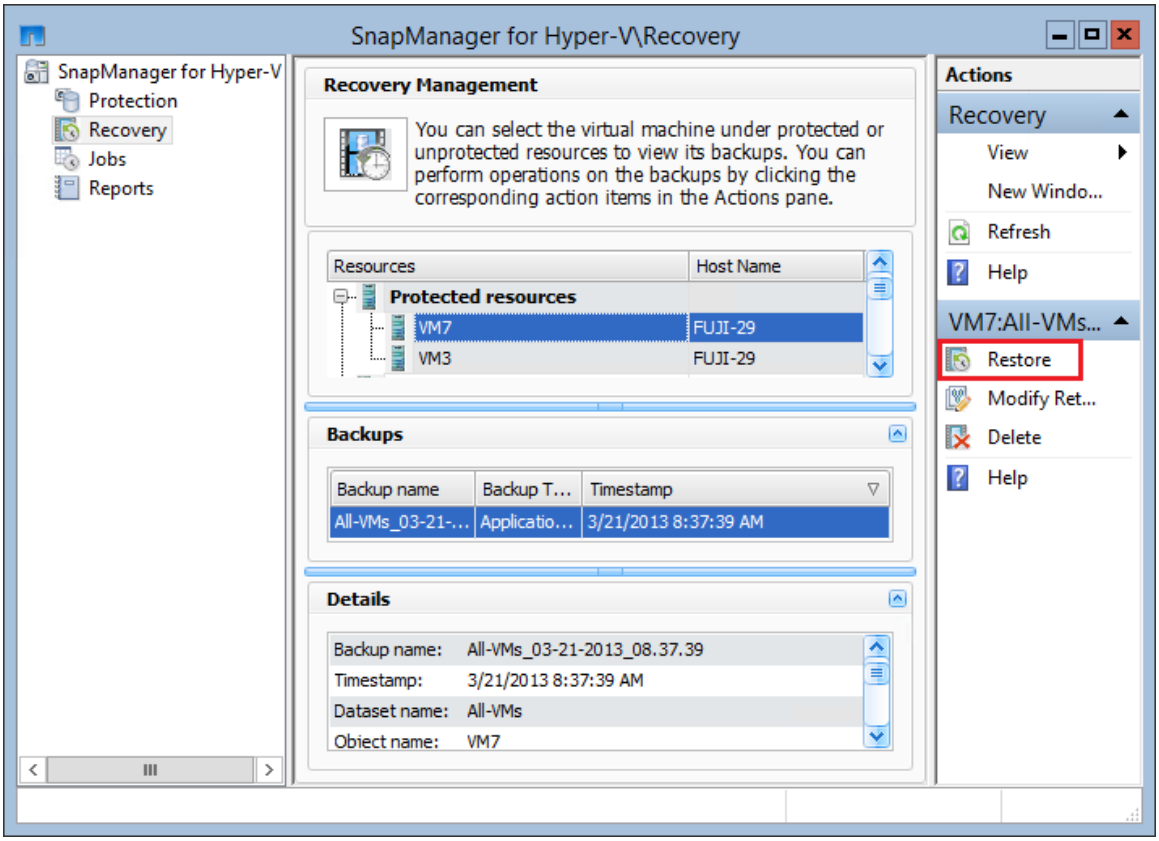


Restore options can be specified as seen in Figure 11.

Figure 11) Restore options.

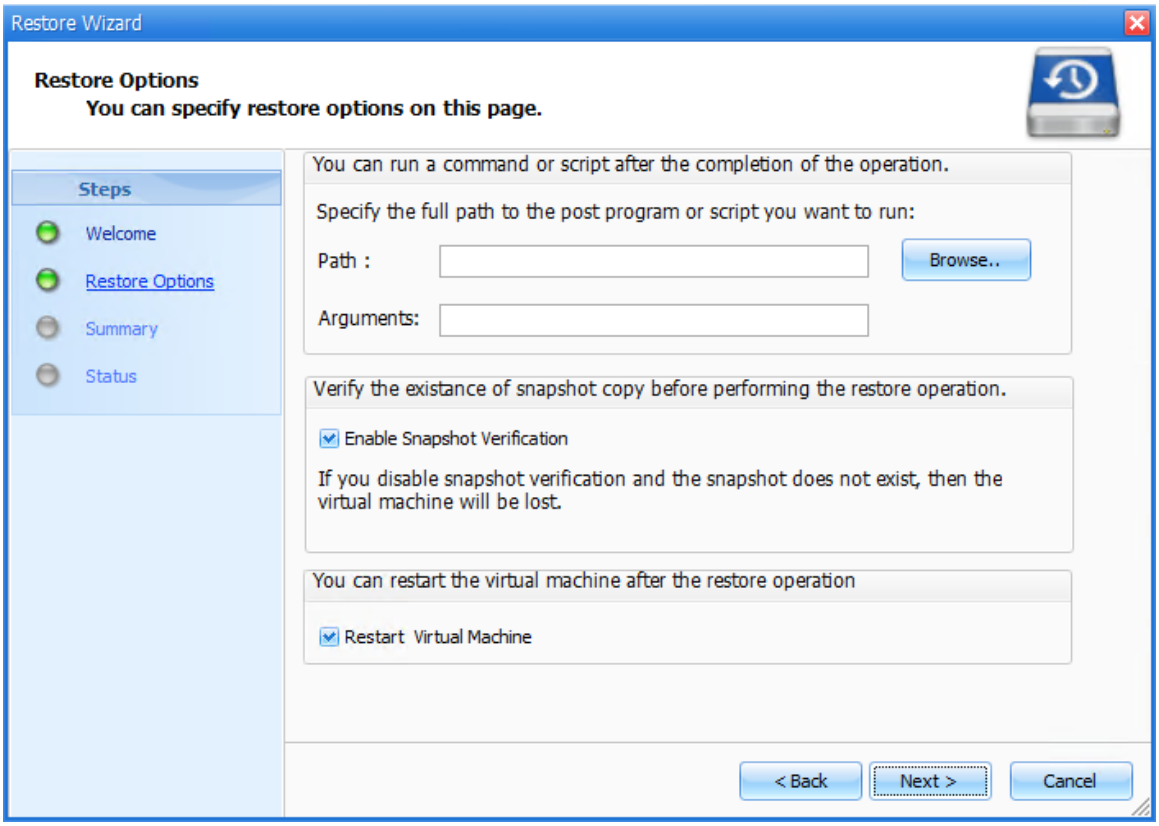

## References

The following references were used in this technical report:

- Windows Server 2012 Overview
  http://www.microsoft.com/en-us/server-cloud/windows-server/overview.aspx
- What's New in Windows Server 2012
  http://technet.microsoft.com/en-us/library/hh831769.aspx
- Deploy Hyper-V over SMB
  http://technet.microsoft.com/en-us/library/jj134187.aspx
- Protect Data on Remote SMB File Share Using VSS
  http://technet.microsoft.com/en-us/library/jj612865.aspx
- Getting Ready for Windows Server 8 Part II – Hyper-V over SMB
  https://communities.netapp.com/community/netapp-blogs/msenviro/blog/2011/09/22/getting-ready-for-windows-server-8-part-ii-hyper-v-over-smb

## Version History

| Version | Date | Document Version History |
|---|---|---|
| Version 1.0 | March 2013 | Initial version, Pavel Lobanov, John Reed |

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

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®


www.netapp.com