



Technical Report

# SnapDrive 6.5 for Windows for Clustered Data ONTAP: Best Practices Guide

Santhosh Harihara Rao, NetApp  
January 2013 | TR-4000

## Abstract

This technical report provides guidelines and best practices for integrated architecture and implementations of Microsoft® Hyper-V™ with NetApp® storage solutions. The NetApp technologies discussed in this technical report are important for achieving an integrated storage solution that is cost effective, operationally efficient, flexible, and environment friendly.

## TABLE OF CONTENTS

<b>1</b>	<b>Overview</b> .....	<b>5</b>
1.1	SnapDrive for Windows Key Features .....	5
<b>2</b>	<b>Scope</b> .....	<b>5</b>
<b>3</b>	<b>Introduction to Clustered Data ONTAP 8.1</b> .....	<b>6</b>
3.1	Characteristics of Clustered Data ONTAP .....	8
<b>4</b>	<b>Achieving High Availability</b> .....	<b>10</b>
<b>5</b>	<b>Sequence of Deployment on the Storage System</b> .....	<b>11</b>
<b>6</b>	<b>Licensing Requirements for SnapDrive 6.5</b> .....	<b>12</b>
<b>7</b>	<b>Deployment Steps on the Windows 2008 R2 Host System</b> .....	<b>12</b>
<b>8</b>	<b>Software Versions Supported in Data ONTAP 7-Mode and Clustered Data ONTAP for Windows Environments</b> .....	<b>14</b>
<b>9</b>	<b>Protocols Supported in SnapDrive for Windows</b> .....	<b>14</b>
<b>10</b>	<b>Upgrading SnapDrive for Windows</b> .....	<b>14</b>
<b>11</b>	<b>Data ONTAP DSM 3.5</b> .....	<b>14</b>
<b>12</b>	<b>ALUA</b> .....	<b>17</b>
<b>13</b>	<b>LUN Management in SnapDrive 6.4 for Windows (or Later)</b> .....	<b>17</b>
13.1	LUN Management During Volume Moves.....	18
<b>14</b>	<b>Snapshot Management</b> .....	<b>19</b>
14.1	Snapshot Scheduling .....	19
14.2	Storage Failover.....	19
<b>15</b>	<b>VMware Support</b> .....	<b>20</b>
<b>16</b>	<b>DataFabric Manager and Role-Based Access Control (RBAC) Support</b> .....	<b>20</b>
<b>17</b>	<b>AutoSupport</b> .....	<b>20</b>
<b>18</b>	<b>Data Protection</b> .....	<b>21</b>
<b>19</b>	<b>SnapManager Dependency on SnapDrive</b> .....	<b>21</b>
<b>20</b>	<b>Data Migration from 7-Mode to Clustered Data ONTAP</b> .....	<b>23</b>
<b>21</b>	<b>Windows Server 2012 Support</b> .....	<b>24</b>

21.1 Introduction .....	24
21.2 Prerequisites .....	24
21.3 Feature Overview.....	24
21.4 Asymmetric Clustering .....	25
21.5 BitLocker Encryption .....	25
21.6 New Virtual Hard Disk Format.....	25
21.7 Hyper-V Virtual Machine Live Migration.....	25
21.8 Hyper-V VM Storage Live Migration.....	25
21.9 Windows Server 2012 Features Not Supported from SnapDrive 6.5 When Connected to NetApp Storage Systems Running in Clustered Data ONTAP .....	26
21.10 Windows Server 2012 Virtual Machine Support for ESX Environments .....	26
<b>22 Performance Tuning.....</b>	<b>26</b>
<b>Appendix.....</b>	<b>26</b>
New Vserver Setup.....	26
New Vserver: LIF Setup .....	27
New Vserver: Routing-Group Setup .....	27
New Vserver: SnapDrive User and Roles.....	28
New Vserver: Create and Start SAN Services.....	29
Set Up with an Existing Vserver .....	30
SnapDrive and Windows Server 2012 Behavior When a New Node Is Added to the Cluster .....	33
<b>Acknowledgement .....</b>	<b>34</b>
<b>Version History .....</b>	<b>35</b>

## LIST OF TABLES

Table 1) Clustered Data ONTAP terminology.....	7
Table 2) Redundancy achieved in each layer.....	10
Table 3) Supported software versions.....	14
Table 4) SnapDrive for Windows.....	14
Table 5) Hotfixes for different operating systems.....	16
Table 6) Recommended hotfixes.....	16

## LIST OF FIGURES

Figure 1) Architecture of clustered Data ONTAP 8.1.....	6
Figure 2) Typical deployment of clustered Data ONTAP 8.1.....	7
Figure 3) Scalability and tiered storage.....	9

Figure 4) Multi-tenancy.....	10
Figure 5) Sample Vserver.....	11
Figure 6) Typical Windows 2008 R2 deployment. ....	13
Figure 7) Option to install Windows Host Utilities as part of DSM 3.5 installation. ....	15
Figure 8) ALUA.....	17
Figure 9) Transport protocol settings page showing C-server and Vserver details. ....	18
Figure 10) Storage failover.....	20
Figure 11) SQL Server virtual machine hosted on clustered Data ONTAP 8.1. ....	23

## 1 Overview

NetApp SnapDrive® for Windows® (SDW) helps you to automate 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 and Windows Server 2012. It can also be used to provision storage for Windows virtual machines hosted on ESX® hypervisors.

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

### 1.1 SnapDrive for Windows Key Features

SnapDrive for Windows:

- Enables online storage configuration, LUN expansion, and streamlined management
- Integrates Data ONTAP® Snapshot™ technology with Microsoft VSS framework, which enables creating point-in-time images of data stored on LUNs
- Works in conjunction with SnapMirror® software to facilitate disaster recovery from either asynchronously or synchronously mirrored destination volumes
- Enables management of SnapDrive on multiple hosts
- Enables support on Microsoft cluster configurations
- Enables iSCSI session management
- Supports iSCSI initiators based on ESX, LUN migration with VMware® vMotion®, and RDM LUNs on virtual machines (VMs) hosted by ESX servers

## 2 Scope

This document focuses on best practices for SnapDrive support for clustered Data ONTAP 8.1 or later. All the features of earlier releases of SnapDrive for Windows are part of this release. This release continues to support management of Data ONTAP 7-Mode storage systems. For best practices for SnapDrive for Windows for 7-Mode systems, refer to [TR-4001: SnapDrive 6.5 for Windows for Data ONTAP 7-Mode: Best Practices Guide](#).

### 3 Introduction to Clustered Data ONTAP 8.1

SnapDrive 6.4 for Windows was released primarily to support NetApp clustered Data ONTAP. This section introduces the clustered Data ONTAP architecture and some key concepts of clustered Data ONTAP in the context of SnapDrive.

Figure 1) Architecture of clustered Data ONTAP 8.1.

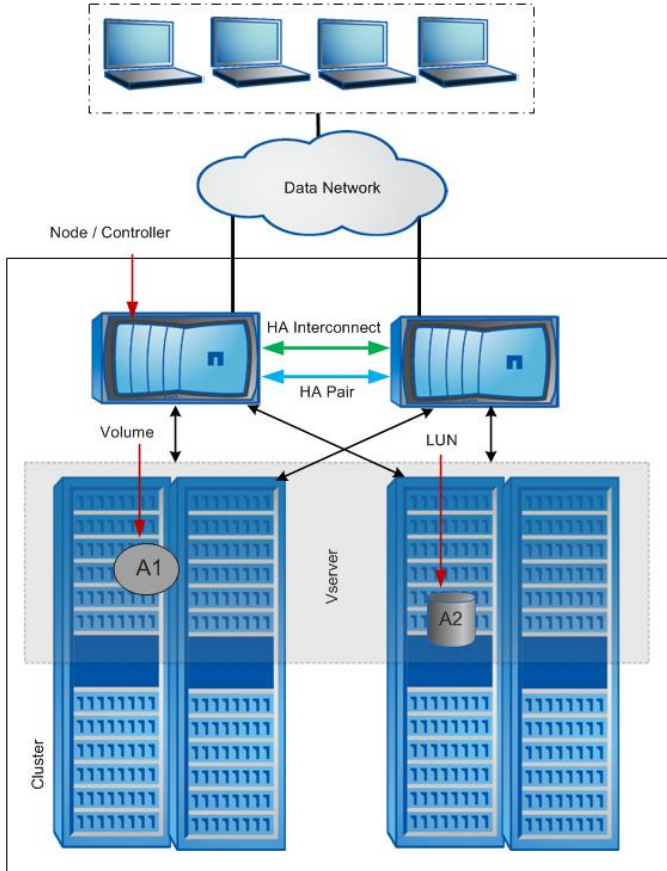
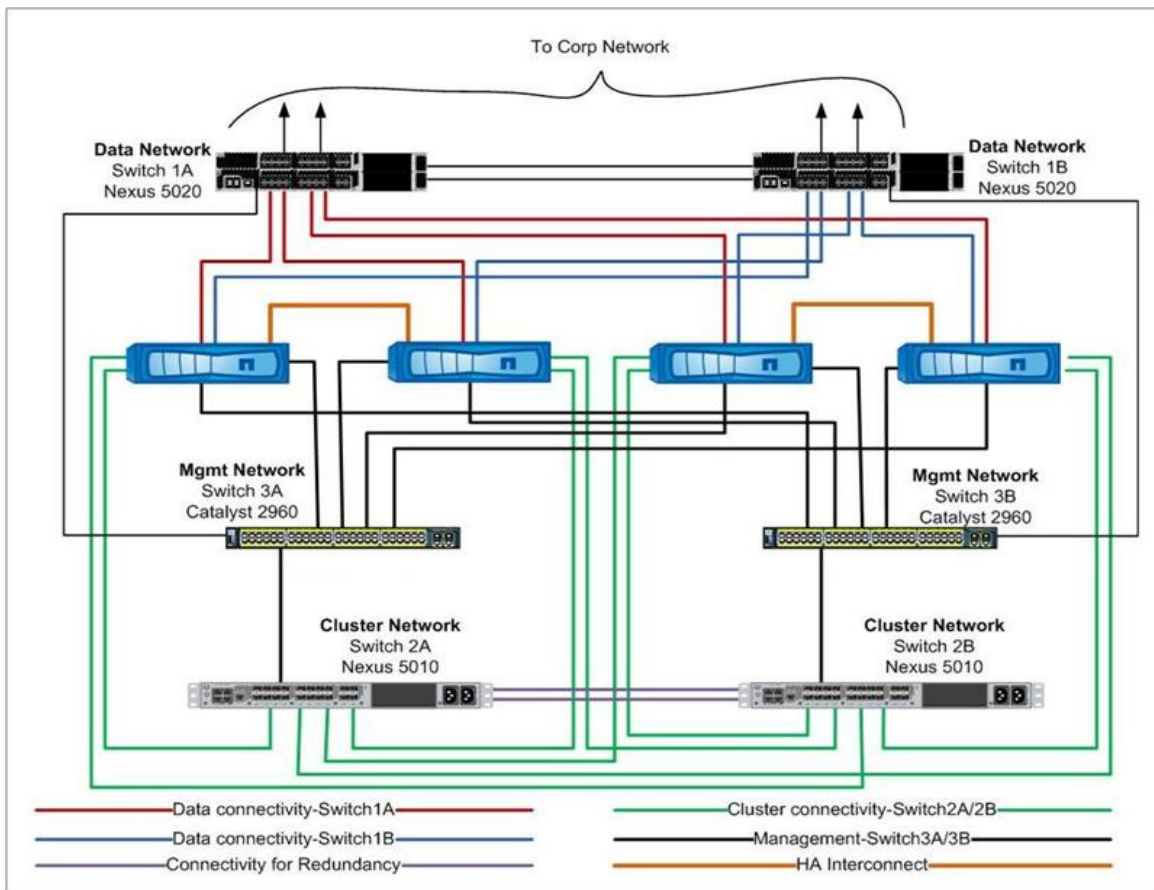


Figure 2) Typical deployment of clustered Data ONTAP 8.1.



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

Table 1) Clustered Data ONTAP terminology.

Term	Description
Cluster	<ul style="list-style-type: none"> <li>In clustered Data ONTAP, a group of connected nodes (storage systems) that share a global namespace and can be used for management as a single virtual server or multiple virtual servers, offering performance, reliability, and scalability benefits.</li> <li>In the Data ONTAP 8.1 release family and earlier releases, a pair of storage systems (sometimes called nodes) configured to serve data if one of the two systems stops functioning.</li> </ul>
Node	One storage controller equals one node; a high-availability (HA) pair equals two nodes.
HA (high availability)	In Data ONTAP, the recovery capability provided by a pair of nodes (storage systems) called an HA pair that are configured to serve data if one of the two nodes stops functioning.
HA pair	In Data ONTAP, a pair of nodes (storage systems) configured to serve data for each other if one of the two nodes stops functioning.
Vserver (virtual server)	<p>A secure virtual storage server that supports multiple protocols and unified storage. A Vserver contains data volumes and one or more logical network interfaces (LIFs) through which it serves data to the clients.</p> <p>A Vserver securely isolates the shared virtualized data storage and network, and it appears to the clients as a single dedicated server. Each Vserver has a separate administrator authentication domain and can be managed independently by a Vserver administrator.</p>

Term	Description
	A Vserver represents a single file system and a unique global namespace. A global namespace enables the NAS clients to access data without specifying the physical location of the data. Global namespace also enables the cluster and Vserver administrators to manage distributed data storage as a single file system.
Admin Vserver	This Vserver type is created when a clustered Data ONTAP system is initialized. This Vserver is responsible only for the administration of the cluster. In SnapDrive, this Vserver interacts with Data ONTAP only for license information and SnapMirror commands, which are both cluster-level concepts.
Cluster Vserver	This Vserver type is created by users for hosting user data, either NAS or SAN. All data objects live here, from data LIFs to LUNs to volumes.
LIF	A logical network interface that represents a network access point to a node. LIFs currently correspond to IP addresses, but they can be implemented by any interconnect. A LIF is generally bound to a physical network port—that is, an Ethernet port. LIFs can fail over to other physical ports (potentially on other nodes) based on policies interpreted by the LIF manager.
Vserver management LIF	A LIF created specifically for managing a Vserver. Although a Vserver management role does not currently exist, this type of LIF allows operations such as Data ONTAP API calls from SnapDrive products to take place. This LIF has a data role and a management firewall policy. It does not host any data protocols (data-protocols set to none).
Routing group	A container for routes, including the default route. In 7-Mode, this was known as a “routing table,” but because clustered Data ONTAP allows multiple Vservers to have their own default routes, NetApp recommends using routing groups to contain them. A default route must be added to the Vserver’s routing groups if routing is necessary for a system or cluster to communicate with hosts.
Cluster interconnect	This is a 10GbE network connection for data communication across nodes.


### 3.1 Characteristics of Clustered Data ONTAP

#### Transparent Data Mobility

With clustered Data ONTAP, users can move volumes from one aggregate to another within a Vserver. This move is nondisruptive and facilitates load balancing, performance optimization, servicing, and upgrades. As a result, clustered Data ONTAP is an ideal candidate for public and private cloud deployments.



Figure 3) Scalability and tiered storage.



## Scalability

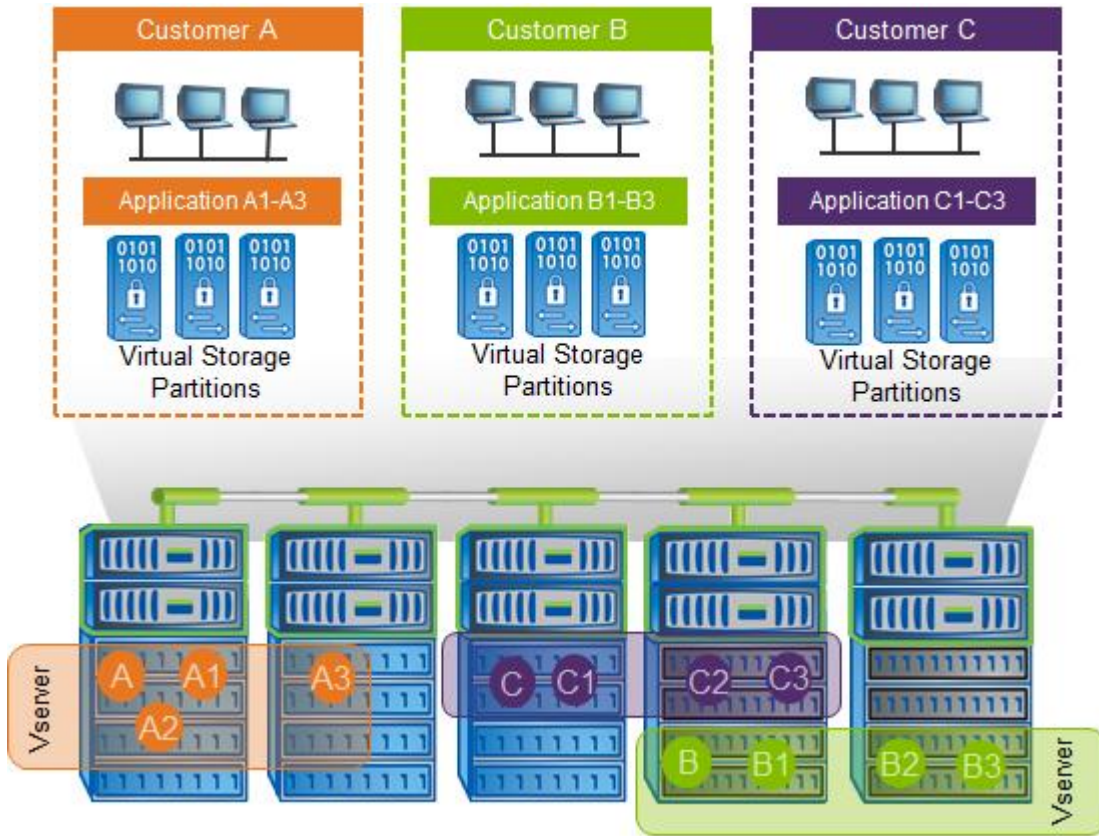
Users can start with a minimal configuration (two-node cluster system) and gradually expand, nondisruptively, as business grows. Users can simultaneously rebalance and fine-tune the capacity of different workloads. In addition, they can group high-performing drives for selected workloads, providing tiered storage options.

## Multi-Tenancy

Multi-tenancy is a critical requirement for cloud deployments. It is imperative that different customers (in the case of public clouds) and different business units (in the case of private clouds) are securely isolated at the compute, network, and storage layers. Clustered Data ONTAP facilitates multi-tenancy at the storage layer by segregating storage entities such as aggregates, LIFs, LUNs, and volumes and containing them in a Vserver. Because each Vserver operates in its own namespace, each unit or customer mapped to this Vserver is completely isolated.

Also, each Vserver supports role-based access control, and specific protocols such as NFS, CIFS, iSCSI, FC, and FCoE can be assigned to it.

Figure 4) Multi-tenancy.



## 4 Achieving High Availability

To create a highly available and redundant environment, make sure that redundancy is addressed in the layers described in Table 2.

Table 2) Redundancy achieved in each layer.

Layer	Description
Sessions from the host to the storage subsystem	Windows multipath I/O (MPIO), Data ONTAP device-specific module (DSM) 3.5
Host Ethernet port	NIC teaming (supported systems only)
Server	Windows failover cluster manager
Ethernet	Two Ethernet switches
FC	Two fabric switches
Storage systems	Two-node clustered Data ONTAP system (minimum)
LIFs in the storage systems	Multiple ports mapped to the same LIF

## 5 Sequence of Deployment on the Storage System

1. Set up the cluster environment. Refer to the Data ONTAP 8.1 Installation and Administration Guide, available on the NetApp [Support](#) site.
2. Create an aggregate.
3. Create a Vserver.
4. Create iSCSI service or FC service to set up an iSCSI or FC target node.

Figure 5) Sample Vserver.

```
APPLCL:> vservershow -vservershow infraserver
Vserver: infraserver
Vserver Type: cluster
Vserver UUID: d3aa16e2-97f6-11e0-bbc3-123478563412
Root Volume: infraroot
Aggregate: infraggr
Name Service Switch: file
Name Mapping Switch: ldap
NIS Domain: -
Root Volume Security Style: ntfs
LDAP Client: -
Language: C
Snapshot Policy: default
Comment:
Anti-Virus On-Access Policy: default
Quota Policy: default
List of Aggregates Assigned: -
Limit on Maximum Number of Volumes allowed: unlimited
Vserver Admin State: running
Allowed Protocols: nfs, cifs, fcp, iscsi
Disallowed Protocols: -
```

5. Configure the network for the Vserver:
  - a. The data LIFs, which enable Vservers to serve data to the clients (iSCSI and FCP)
  - b. The management LIF, which allows SnapDrive to communicate with the other LIFs to serve data
6. Create data volumes of the required size. SnapDrive uses these volumes to create and manage LUNs.
7. For data protection within the cluster, perform the following additional steps:
  - a. Create a volume in the Vserver of the secondary Vserver. Make sure that the property of that volume is of the DP type.
  - b. Establish a SnapMirror relationship between the primary and the secondary by accessing the secondary system.
8. For intercluster SnapMirror replication, make sure that at least one intercluster management LIF is present in each node on both primary and secondary storage systems. For additional information, see subsequent sections.

For more information about data protection, refer to:

- [Data ONTAP 8.1 Cluster-Mode Data Protection Guide](#)
- [SnapMirror FAQ](#)

**Note:** SnapDrive does not support single sign-on when connected to clustered Data ONTAP. Users must configure user credentials for all cluster Vserver and Vserver management endpoints from the storage system that will be accessed by SDW.

## Best Practices

In the storage system, NetApp recommends at least four LIFs per Vserver:

- Two data LIFs
- One management LIF
- One intercluster LIF (for intercluster replication)

## 6 Licensing Requirements for SnapDrive 6.5

To enable all of the features, NetApp requires that the following licenses be present in the clustered Data ONTAP storage system:

- FlexClone<sup>®</sup>
- iSCSI
- FCP
- SnapRestore<sup>®</sup>
- SnapMirror (optional)
- SnapManager<sup>®</sup> suite (if the user is using storage system–based licensing)
- MPIO license on the storage system

**Note:** All of the SnapManager products for Windows environments—SnapManager for Hyper-V, SnapManager for SQL Server<sup>®</sup>, SnapManager for SharePoint<sup>®</sup>, and SnapManager for Exchange—communicate directly with SDW for the host license check.

## 7 Deployment Steps on the Windows 2008 R2 Host System

Here are the deployment steps for a Windows 2008 R2 server system.


1. Install the following hotfixes:
  - a. Q2522766 (for DSM 3.5)
  - b. Q2528357 (for DSM 3.5)
  - c. KB2494016 (virtual machines crash when CSV goes into redirected mode)
  - d. KB2520235 (stop error when more than 200 hard disks are presented to a WSFC)
  - e. KB2531907 (SCSI device vital product data [VPD] test runs on online disks, causing data corruption)
2. Install Data ONTAP DSM 3.5.
3. Install SnapDrive 6.4 for Windows or later.

In SnapDrive 6.5 for Windows, add the management LIF (created while creating the Vserver) in Add Storage Systems under the Transport Protocol option. All the data is managed through the management LIFs that belong to this Vserver in SnapDrive. Again, add the IP of the management LIF in the preferred storage system IP addresses.

4. For managing clustered Data ONTAP LUNs, select HTTP or HTTPS as the transport protocol. The RPC protocol is not supported.
5. In SnapDrive for Windows or later, create the LUN in the selected Vserver.
6. To establish a SnapMirror relationship, add the Vserver credentials of the source and destination volumes in the primary host.

Figure 6 shows the deployment of cluster shared volumes (CSVs) hosting guest VMs and application data using SDW and storage systems based on clustered Data ONTAP.

Figure 6) Typical Windows 2008 R2 deployment.



## 8 Software Versions Supported in Data ONTAP 7-Mode and Clustered Data ONTAP for Windows Environments

Table 3) Supported software versions.

7-Mode	7-Mode and Clustered Data ONTAP
Data ONTAP 8.0 or earlier	Data ONTAP 8.1 or later
DSM 3.4 or earlier	DSM 3.5 or later
SnapDrive 6.3 or earlier for Windows	SnapDrive 6.4 for Windows or later
Windows Host Utility Kit 5.0	No longer required for DSM; the DSM 3.5 installer includes the MBRAAlign tool and the LinuxGuestConfig.iso in the installation package. If the user has Microsoft DSM (MSDSM) installed, then Windows Host Utilities 6.0 must be installed.
SnapManager for Hyper-V 1.0 P1	SnapManager for Hyper-V 1.0 P1 or later
SnapManager for SQL Server 5.1	SnapManager for SQL Server 5.2 or later
SnapManager for SharePoint 6.0	SnapManager for SharePoint 6.1 or later
SnapManager for Exchange 6.0.2	SnapManager for Exchange 6.0.2 or later
ESXi™ 5.0 and previous releases	ESXi 5.0 and previous releases

## 9 Protocols Supported in SnapDrive for Windows

Table 4) SnapDrive for Windows.

Parameter	7-Mode Storage Systems	Clustered Data ONTAP Storage Systems
Storage protocol	SAN over iSCSI or FC or FCoE For ESX environments: VMDK over NFS and VMDK over VMFS	SAN over iSCSI or FC or FCoE For ESX environments: VMDK over NFS and VMDK over VMFS
Network protocol	IPv4 and IPv6	IPv4 only
Transport protocol	RPC, HTTP, and HTTPS	HTTP and HTTPS

## 10 Upgrading SnapDrive for Windows

If the user is adding a clustered Data ONTAP storage system to the same host that is using 7-Mode systems, then before the SnapDrive upgrade is performed, to enable multipathing, the DSM on the host system must be upgraded to allocate storage from clustered Data ONTAP storage systems. Section 11, “Data ONTAP DSM 3.5,” describes this aspect.

## 11 Data ONTAP DSM 3.5

Microsoft MPIO is a protocol-independent feature that supports multiple data paths to a storage device with iSCSI, Fibre Channel, or SAS. Providing multiple paths that can handle failover increases the



availability from a host to the storage system. Windows 2008 R2 x 64 servers include support for Microsoft MPIO.

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

DSM.3.5 is released to support host connections to clustered Data ONTAP storage systems. In addition, by upgrading to DSM 3.5 from DSM 3.4, users can manage both 7-Mode and clustered Data ONTAP storage systems.

The DSM 3.5 installer includes all settings and files required for the driver. This eliminates the requirement of Windows Host Utilities with Data ONTAP DSM 3.5. The installer includes the MBRAlign tool and the LinuxGuestConfig.iso in the installation package. The installer displays a dialog box prompting the user to install the Hyper-V Guest Utilities; the default value is No. Selecting Yes installs the MBRAlign tool and LinuxGuestConfig.iso in the INSTALLDIR directory (default is %systemdrive%\Program Files\NetApp\MPIO).

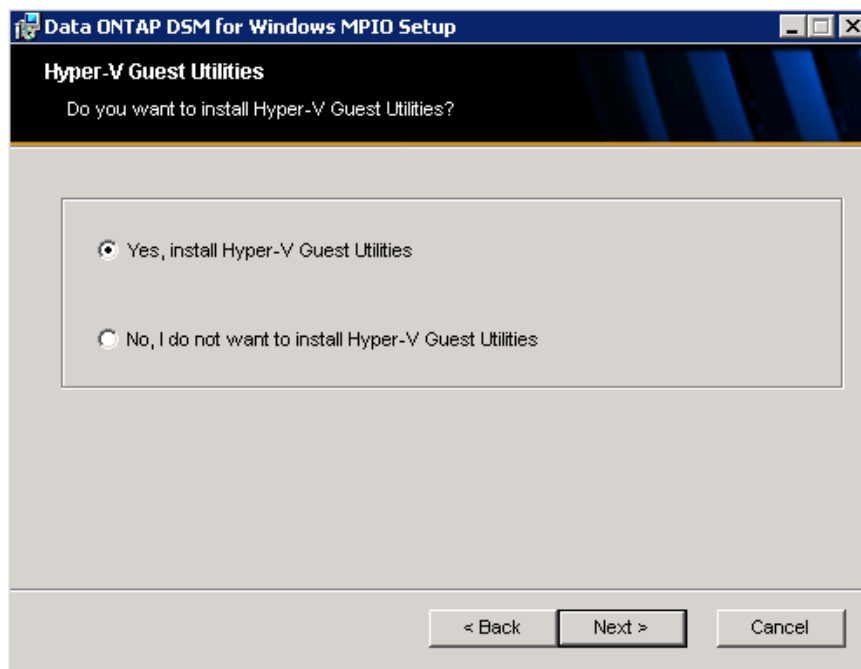
#### Best Practice

For a highly available connection to the storage system, NetApp requires installing the supported version of multipathing software such as Data ONTAP DSM for Windows MPIO.

**Note:** The Windows Host Utilities Kit, which included time-out value settings for Hyper-V VHDs, is no longer needed in SnapDrive for Windows 6.5 and later. This kit is already bundled as part of DSM 3.5.

Note: DSM 3.5 does not support Windows Server 2012. In Windows Server 2012, the user must use MSDSM.

Figure 7) Option to install Windows Host Utilities as part of DSM 3.5 installation.



## Best Practices

Follow these steps to upgrade the host system to DSM 3.5 and SnapDrive 6.4 for Windows(or later) on a Windows 2008 R2 SP1 host:

1. Uninstall the Windows Host Utilities Kit present on the Windows 2008 host system.
2. Reboot the host system.
3. Uninstall DSM 3.4 and reboot the host system.
4. Install Q2522766 and Q2528357 hotfixes from Microsoft.
5. Install DSM 3.5.
6. Reboot the host system.
7. Upgrade from SnapDrive 6.3 to SnapDrive 6.4 or later.

**Note:** If the user has MSDSM installed on the host and has not installed Data ONTAP DSM 3.5, then Windows Host Utilities 6.0 must be installed to support clustered Data ONTAP LUNs.

Table 5 contains the complete list of hotfixes for different operating systems.

Table 5) Hotfixes for different operating systems.

Operating System	Hotfix List
Windows 2003 R2 SP2	<a href="#">945119</a> Stop error that is related to the Storport.sys driver on a Windows Server 2003-based computer: "0x000000D1 DRIVER_IRQL_NOT_LESS_OR_EQUAL." <a href="#">982109</a> Access to an MPIO-controlled storage device fails after you disconnect all data cables in Windows Server 2003.
Windows 2008 SP2	<a href="#">968675</a> The storport.sys driver leaks nonpaged pool memory when volumes are scanned, mounted, or dismounted in Windows Server 2008 and Windows Vista®. <a href="#">2522766</a> The MPIO driver fails over all paths incorrectly when a transient single failure occurs in Windows Server 2008 or in Windows Server 2008 R2.
Windows 2008 R2	<a href="#">2522766</a> The MPIO driver fails over all paths incorrectly when a transient single failure occurs in Windows Server 2008 or in Windows Server 2008 R2. <a href="#">2528357</a> Nonpaged pool leaks when you disable and enable some storage controllers in Windows 7 or in Windows Server 2008 R2. <a href="#">979711</a> Stop error message on a computer that is running Windows 7 or Windows Server 2008 R2 and that has iSCSI storage: "0x0000000A."
Windows 2008 R2 SP1	<a href="#">2522766</a> The MPIO driver fails over all paths incorrectly when a transient single failure occurs in Windows Server 2008 or in Windows Server 2008 R2. <a href="#">2528357</a> Nonpaged pool leaks when you disable and enable some storage controllers in Windows 7 or in Windows Server 2008 R2.

Table 6 lists the recommended hotfixes (not enforced by the installer but recommended for host clustered configurations).

Table 6) Recommended hotfixes.

Operating System	Hotfix List
Windows 2008 R2	<a href="#">2520235</a> "0x0000009E" stop error when you add an additional storage disk to a failover cluster in Windows Server 2008 R2.
Windows 2008 R2 SP1	<a href="#">2520235</a> "0x0000009E" stop error when you add an additional storage disk to a failover cluster in Windows Server 2008 R2. <a href="#">2531907</a> Validate SCSI device VPD test fails after you install Windows Server 2008 R2 SP1.



## 12 ALUA

Asymmetric Logical Unit Access (ALUA) support is available on all clustered Data ONTAP configurations with Data ONTAP DSM as well as with MSDSM.

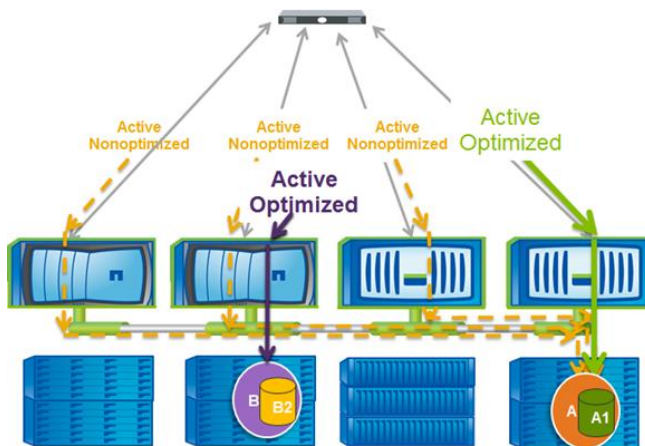
Hosts use ALUA to determine the state of a specified path:

- Active/optimized: The LIF is configured on a port on a controller that also “owns” the aggregate on which the LUN is provisioned.
- Active/unoptimized: The LIF is configured on a port on a controller that does not own the aggregate on which the LUN is provisioned.

Clustered Data ONTAP by default enables ALUA. When creating a shared disk, take into account the following considerations with respect to the host server OS that is used:

- For Windows 2003, it is mandatory to install Data ONTAP DSM 3.5 for clustered Data ONTAP.
- Windows 2008, Windows 2008 SP2, and later should be used for clustered Data ONTAP shared LUNs.

Figure 8) ALUA.



## 13 LUN Management in SnapDrive 6.4 for Windows (or Later)

With clustered Data ONTAP 8.1, users can isolate a single workload within a Vserver. A Vserver functions as a logic container that has its own namespace. After a user creates the Vserver, root volume, and data volume, SnapDrive for Windows can be installed on the host server to perform LUN management tasks such as create, expand, destroy, connect, disconnect, and so on.

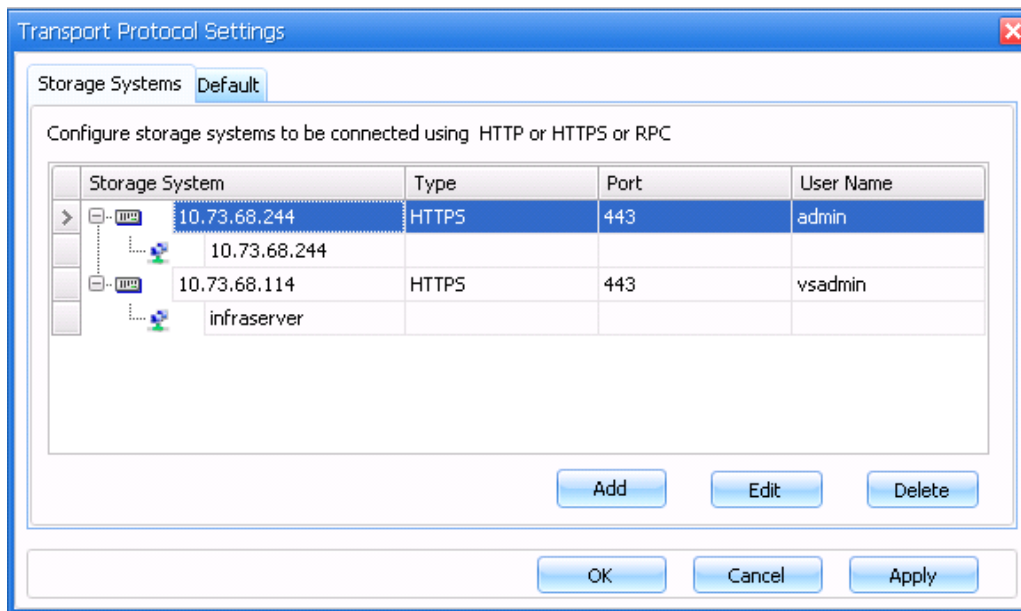
Before performing any LUN management task, users must add the management LIF that was created during the Vserver creation. This LIF communicates internally with other data LIFs such as iSCSI or FC LIFs; it also performs tasks.

### Best Practice

At least one data LIF and one management LIF per Vserver are required for SnapDrive to perform LUN management tasks.

Before performing any operation, the Vserver management LIF in SnapDrive > Transport Protocols Settings > Storage Systems must be entered. In addition, if the user needs to perform activities related to SnapMirror and ASUP™ and Snapshot schedule management, cluster server admin credentials must be entered.

Figure 9) Transport protocol settings page showing C-server and Vserver details.



### 13.1 LUN Management During Volume Moves

Clustered Data ONTAP 8.1 supports volume moves from one aggregate to another aggregate anywhere within the cluster. Volume moves are performed when cluster expansion is needed or when the business priority changes or as part of a hardware lifecycle management task.

A volume is moved by using a single administrator command (from the CLI or System Manager 2.0). Volume moves are nondisruptive and provide uninterrupted access to the hosts that are connected to the storage using any protocol—iSCSI, NFS, CIFS, FC, or FCoE—during the volume move.

This means that users can continue to access data while a volume move is being performed in the background. All sessions from SnapDrive to the storage systems continue to exist.

**Note:** SnapDrive cannot detect the volumes that are currently in a moving state. All volumes are displayed in the Create/Connect wizard, irrespective of their states.

#### Best Practices

- The cutover window defined for a SAN volume should not exceed the expected timeout value on the host side. During the cutover phase of the volume move, all I/O access is queued, and requests are blocked to the source volume. SnapDrive sets a timeout value of 120 seconds on the host during the volume move. In addition, when a SAN volume is moved, ALUA is used for optimized access to the volume. NetApp recommends performing volume moves during nonpeak hours.
- Each node must have a data LIF for optimized access to the volume.

**Note:** SnapDrive does not support file-based thin provisioning in clustered Data ONTAP.

## 14 Snapshot Management

Similar to LUN management, Snapshot copies are managed within a Vserver. SnapDrive initiates a request for Snapshot copies on the node that is hosting the Vserver.

When SDW initiates a Snapshot copy on a LUN present in a 7-Mode system, a LUN clone operation is initiated. This process uses NetApp FlexClone technology in the back end. In addition, for restoring a LUN, the LUN clone split restore operation is initiated.

In the case of clustered Data ONTAP systems, LUN Snapshot copy creation and restoration are initiated by using SIS cloning. No impact is observed on the SnapDrive application because this activity is performed on the storage system.

**Note:** No separate license is required to enable SIS cloning on clustered Data ONTAP systems.

### 14.1 Snapshot Scheduling

When SnapDrive manages LUNs in a 7-Mode system, SDW turns off the Snapshot copy schedule whenever a LUN is provisioned by SDW on a volume. The intention is to manage Snapshot copies on the volume if it contains SDW-managed LUNs to save space.

In the case of clustered Data ONTAP systems, the Snapshot copy scheduling can be performed only if the C-server credentials are provided, because it is the responsibility of the C-server administrator to manage space allocation on the Vservers.

#### Best Practices


- Depending on the level of security required, the administrator can decide whether to provide C-server credentials on the host side to enable Snapshot copy scheduling. Otherwise, the C-server administrator (cluster administrator) should reserve adequate space for Snapshot copies while creating the Vservers and aggregates for volume provisioning. This is necessary because the Vserver administrator does not have sufficient privileges to add storage space to the aggregate.
- Because the SnapManager suite of products (SnapManager for SQL Server, SnapManager for SharePoint, SnapManager for Exchange, and SnapManager for Hyper-V) uses SnapDrive for application-consistent Snapshot copies, NetApp recommends making sure that there are minimal overlaps when these application-specific Snapshot copies are initiated through their respective products. This reduces the performance overhead on the C-server.

### 14.2 Storage Failover

To accommodate storage failover, two storage controllers (nodes) in the same cluster are connected together as an SFO pair, called an “active-active” pair. Each node of the pair is a fully functioning node in the cluster (hence the term “active-active”). Clusters can be heterogeneous (in terms of hardware and Data ONTAP versions), but an SFO pair must be the same controller model.

**Note:** During a storage failover, any activity related to SnapDrive and SnapManager will fail. All operations continue to work after the failover is complete.

Figure 10) Storage failover.



## 15 VMware Support

For clustered Data ONTAP, SDW will continue to support guest virtual machines residing on ESX and ESXi hosts. For best practices for 7-Mode support for ESX on SnapDrive, refer to [TR-4001: SnapDrive 6.5 for Windows for Data ONTAP 7-Mode: Best Practices Guide](#).

**Note:** Clustered Data ONTAP requires VSC 4.1 and SnapDrive 6.4.2 for Windows or later to support VMDKs on NFS and VMFS datastores.

## 16 DataFabric Manager and Role-Based Access Control (RBAC) Support

For Data ONTAP 7-Mode storage systems, to enforce RBAC rules, SnapDrive leverages the user's RBAC configuration file (AccessControl.xml), which is located in the root volume of the storage system. This is called file-based RBAC for SnapDrive. If RBAC is configured in DataFabric® Manager (DFM), SnapDrive communicates with DFM and enforces the RBAC rules that have been configured. This is called DFM-based RBAC for SnapDrive. In SnapDrive for Windows, both of these continue to be supported only for 7-Mode systems.

**Note:** SnapDrive for Windows does not support RBAC for clustered Data ONTAP.

For best practices for 7-Mode support for RBAC on SnapDrive, refer to [TR-4001: SnapDrive 6.5 for Windows for Data ONTAP 7-Mode: Best Practices Guide](#).

## 17 AutoSupport

SnapDrive supports AutoSupport™ (ASUP) for both systems based on both 7-Mode and clustered Data ONTAP. SnapDrive LUN and Snapshot copy management events for clustered Data ONTAP can be sent to their respective storage systems. For SDW to pass SMx-specific ASUP data such as SMSQL events, you must have ASUP enabled.

**Note:** To receive SDW ASUP alerts for clustered Data ONTAP, C-server credentials must be provided in the SDW storage settings.

### Best Practice

In a secure multi-tenant environment, it is not advisable to share the C-server credentials on the host/guest VM. In such cases, it is left to the discretion of the customer whether to enable ASUP on the host/guest VM side.

## 18 Data Protection

### SnapMirror

Clustered Data ONTAP 8.1 allows two types of replication:

- Intercluster asynchronous volume replication (replication between volumes hosted on different clusters) for enabling disaster recovery replication to a cluster in a remote site
- Intracluster replication between two Vservers in the same cluster

**Note:** Users cannot establish a SnapMirror relationship between a 7-Mode source volume and a clustered Data ONTAP destination volume. Also, qtree-based SnapMirror is not supported in clustered Data ONTAP.

**Note:** Intercluster replication requires at least one intercluster logical interface (intercluster LIF) per node. The intercluster LIF can be assigned to a data port or to a dedicated intercluster port.

### Best Practice

NetApp recommends having adequate bandwidth over a WAN for the initial transfer.

SnapDrive supports both intercluster and intracluster replication. At any given time, SnapDrive has two Snapshot copies created for SnapMirror transfers.

**Note:** To update SnapMirror in SnapDrive, you must add both the cluster management IP (C-server) credentials and Vserver credentials of the source volume and the destination volume.

### Best Practice

In a secure multi-tenant environment deployed in a clustered Data ONTAP system, it is not advisable to share the C-server credentials on the host for security reasons. In such cases, it is left to the discretion of the customer whether to update SnapMirror using SnapDrive.

### SnapVault

Clustered Data ONTAP does not have SnapVault® software capability; therefore, SDW does not support SnapVault when connected to clustered Data ONTAP systems.

### IPv6

Clustered Data ONTAP does not support IPv6; therefore, SnapDrive does not support IPv6 access to clustered Data ONTAP clusters.

## 19 SnapManager Dependency on SnapDrive

The entire SnapManager suite of products for Windows environments—SnapManager for Hyper-V, SnapManager for SQL Server, SnapManager for SharePoint, and SnapManager for Exchange—communicates directly with SDW for application-consistent Snapshot copies.

## Best Practice

Because the SnapManager suite of products uses SnapDrive for application-consistent Snapshot copies, NetApp recommends making sure that no overlaps occur when these application-specific Snapshot copies are initiated using their respective products. This reduces the performance overhead on the storage systems.


**Note:** To receive application-specific ASUP data, ASUP must be enabled for the SnapManager software, and C-server credentials must be provided in SnapDrive. For example, if you require SMSQL ASUP data, then you must enable ASUP in SMSQL so that SnapDrive receives the data and sends it to the storage system.




**Note:** The SnapManager\_suite license must be provided on all nodes of the clustered Data ONTAP system.

## Best Practices

- Store the applications and the virtual machines that host them in a single Vserver. This facilitates secure isolation and simplifies backup restores.
- Store the applications and the virtual machines that host them in a single Vserver. The volumes in this Vserver can have a SnapMirror relationship with another Vserver containing an identical set of volumes. This simplifies the disaster recovery strategy.

Figure 11) SQL Server virtual machine hosted on clustered Data ONTAP 8.1.



	Volume in Vserver-1 containing CSV LUN that hosts SQL Server VM.
	LUN hosting the database for SQL Server application. This LUN is hosted on another volume within Vserver-1.
<b>Vserver-1</b>	Vserver-1 hosts all the data volumes that are specific to SQL Server (virtual machines and database.)
<b>SnapDrive</b>	SnapDrive for Windows installed on the host server.
	SnapManager for SQL Server and SnapDrive for Windows installed on the Guest VM.

## 20 Data Migration from 7-Mode to Clustered Data ONTAP

Volume moves from 7-Mode systems to clustered Data ONTAP clusters are not supported. Also, establishing a SnapMirror relationship between a 7-Mode source volume and a clustered Data ONTAP destination is not supported.



## Best Practices

To migrate data from a 7-Mode system to a clustered Data ONTAP system, NetApp recommends using the data migration features offered by the application. These methods can be used if the data size to be migrated is small (< 1TB). Here are two use cases:

- **Moving Hyper-V virtual machines.** Mount both clustered Data ONTAP and 7-Mode LUNs (hosting the virtual machines) using SnapDrive. Perform storage migration using SCVMM. If SCVMM is not part of the infrastructure, shut down the virtual machines and copy the entire virtual machine (.vhd and .vmc files) to the new LUN. For more information, refer to <http://technet.microsoft.com/en-us/library/cc708392%28WS.10%29.aspx>.
- **Moving SQL Server databases.** Mount both clustered Data ONTAP and 7-Mode LUNs (hosting the virtual machines) using SnapDrive. Use SQL Server confer wizard for migration. For more information, refer to <http://support.microsoft.com/kb/224071>.

## 21 Windows Server 2012 Support

### 21.1 Introduction

SnapDrive for Windows supports Windows Server 2012. SnapDrive 6.5 for Windows is a prerequisite software that must be installed on Windows Server 2012 to manage NetApp storage systems. SnapDrive 6.5 for Windows will support managing only SAN environments in Windows Server 2012.

**Note:** SnapDrive 6.5 for Windows is a prerequisite for the SnapManager suite of products to function in Windows Server 2012 environments.

### 21.2 Prerequisites

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

### 21.3 Feature Overview

SnapDrive 6.5 for Windows will support all major SAN-based features in Windows Server 2012.

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

## CSV 2.0 Support (CSVFS)

In Windows Server 2012, cluster-shared volumes have undergone significant changes with respect to security, performance, and file system availability for additional cluster workloads. A new clustered file system has been introduced, and this functions as a layer of abstraction above the NTFS file system for the storage volume. As a result, simultaneous read/writes can be performed on the CSV LUN from different nodes. For more details on CSV 2.0, refer to <http://technet.microsoft.com/en-us/library/jj612868.aspx>.

**Note:** A CSV 2.0 volume will have two volume GUIDs:

- **NTFS volume GUID.** When disk is created and partitioned with NTFS and before adding to the CSV.
- **CSV volume GUID.** When disk is added to the cluster-shared volumes.



SnapDrive 6.5 for Windows supports CSVFS.

#### Best Practice

It is recommended in SDW to create a CSV from the node that owns the available cluster storage group. Use the "CLUSTER GROUP" command or "Get-Cluster Group" cmdlet to identify the node that owns the "Available Storage" group before creating a CSV disk.

## Asymmetric Clustering

Asymmetric clustering is a feature by which a user can create a shared disk or cluster-shared volume only among a few nodes in a cluster.

**Note:** SnapDrive 6.5 for Windows does not support this feature.

## BitLocker Encryption

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

SnapDrive 6.5 for Windows will support BitLocker functionality for CSV provisioned through it.

## New Virtual Hard Disk Format

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

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

**Note:** SnapDrive 6.5 for Windows currently cannot create LUNs beyond 16TB due to limitations in Data ONTAP, and hence it is advised to create VHDx for size lesser than 16TB and use other means of provisioning additional storage (for example, pass-through disks, guest iSCSI initiator) using SnapDrive 6.5 for Windows within the VM.

## Hyper-V Virtual Machine Live Migration

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

#### Best Practice

Operations related to SnapDrive within the virtual machine are best avoided during live migration.

## Hyper-V VM Storage Live Migration

This feature in Windows Server 2012 enables migration of virtual machine–related files to a different storage location without the VM having to undergo downtime. SnapDrive 6.5 for Windows supports this feature. It is no longer required to make virtual machine state offline when migrating to a different storage system. Refer to the following link for more information on VM storage live migration:

<http://technet.microsoft.com/pt-br/library/hh831656.aspx>.

### Best Practice

Operations related to SnapDrive during storage live migration are best avoided because they might result in the corruption of the virtual machine.

## 21.4 Windows Server 2012 Features Not Supported from SnapDrive 6.5 When Connected to NetApp Storage Systems Running in Clustered Data ONTAP

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

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

## 21.5 Windows Server 2012 Virtual Machine Support for ESX Environments

ESXi 5.0 U1 and ESXi 5.1 support Windows Server 2012 virtual machines. NetApp Virtual Storage Console (VSC) 4.1 does not support Windows Server 2012.

## 22 Performance Tuning

To improve the performance of the SnapDrive system, follow this best practice.

### Best Practice

Add exclusions for SnapDrive service in the antivirus application agent (NetApp KB ID 2013167 and 2013060) installed on the host server system. Also, configure the Windows firewall to allow SnapDrive services for Windows communications.

## Appendix

### New Vserver Setup

To set up a new Vserver for use with LUNs and SnapDrive, follow these steps.

1. Create a Vserver.

```
::> vservers create -vserver <cluster vserversname> -rootvolume <root volume name> -aggregate <aggr name> -ns-switch <file|nis|ldap> -nm-switch <file|ldap> -rootvolume-security-style <ntfs|unix|mixed>
```

Example:

```
::> vservers create -vserver sdw_vs -rootvolume snapdrive_root -aggregate aggr1 -ns-switch file -nm-switch file -rootvolume-security-style ntfs
```

2. Create a data volume. For size considerations, review the Block Access Management Guide for iSCSI and FC that is specific to your version of clustered Data ONTAP (available on the NetApp [Support](#) site).

```
::> volume create -vserver <cluster vservername> -volume <volname> -aggregate <aggrname> -size <vol size> -security-style <ntfs|unix|mixed>
```

Example:

```
::> volume create -vserver sdw_vs -volume sdw_vol -aggregate aggr1 -size 5g -security-style ntfs
```

## New Vserver: LIF Setup

SnapDrive requires a minimum of two LIFs (one Vserver management, one SAN) for each cluster Vserver, as well as a cluster management LIF for the admin Vserver. SnapDrive communicates with the cluster over the Data ONTAP API to initiate tasks such as LUN create, snap create, LUN clone, and so on. Because ZAPI is an RPC protocol, it must communicate with the system over an Ethernet interface.

**Note:** SnapDrive is not involved in the block-layer transfer of data, initiators, and targets. This is taken care of by the SAN infrastructure.

For more information on LIFs in clustered Data ONTAP, refer to the Network Management Guide that is specific to your release of clustered Data ONTAP (available on the NetApp [Support](#) site).

1. Create LIFs for the Vserver.

## Vserver Management LIF

```
::> net int create -vserver <cluster vservername> -lif <vserver mgmt lif name> -role data -data-protocol none -home-node <home node> -home-port <home port> -address <IP address> -netmask <netmask> -routing-group <routing-group-name-optional>
```

Example:

```
::> net int create -vserver sdw_vs -lif vs_mgmt -role data -data-protocol none -home-node node1 -home-port e0c -address 10.61.65.183 -netmask 255.255.255.0
```

SnapDrive does not require a LIF for NAS. The preceding example shows how to create a LIF to enable NFS and CIFS traffic to the Vserver. A single Vserver can easily support both SAN and NAS traffic, but NetApp recommends not storing both SAN and NAS data on the same volume.

## SAN Data LIFs (Required for Block Data Access)

1. For iSCSI (create one on each node, up to four):

```
::> net int create -vserver <cluster vservername> -lif <sanlif name1> -role data -data-protocol iscsi -home-node <home node1> -home-port <home port1> -address <ip address> -netmask <netmask> -routing-group <routing-group-name-optional>
```

Example:

```
::> net int create -vserver sdw_vs -lif vs_iscsi1 -role data -data-protocol iscsi -home-node node1 -home-port e4a -address 10.61.65.188 -netmask 255.255.255.0
```

2. For FCP (create one on each node, up to four):

```
::> net int create -vserver <cluster vservername> -lif <fcp lif name1> -role data -data-protocol fcp -home-node <home node1> -home-port <home port1>
```

Example:

```
::> net int create -vserver sdw_vs -lif vs_fcpl -role data -data-protocol fcp -home-node node1 -home-port e2
```

## New Vserver: Routing-Group Setup

When creating a LIF, you can specify a routing group. If you do not, a routing group is created automatically when the LIF is created. However, a default route is not created automatically. A default

route must be created for IP traffic to reach the intended router or gateway. Routes added to a routing group must live in the same subnet as specified for the routing group. For more information on routing groups, refer to the Network Management Guide for your specific release of clustered Data ONTAP (available on the NetApp [Support](#) site).

1. Make sure that a default route for the Vserver is created.

```

::> routing-groups route show -vserver <cluster vservname>

::> routing-groups route create -vserver <cluster vservname> -routing-group <routing-group
name> -destination <destination> -gateway <gateway>

```

Example:

```

::> routing-groups route create -vserver sdw_vs -routing-group d10.61.65.0/24 -destination
0.0.0.0/0 -gateway 10.61.65.1

```

2. Verify that the new routing group is correct.

```

::> routing-groups route show -vserver sdw_vs
(network routing-groups route show)
      Routing
Vserver  Group      Destination      Gateway      Metric
-----  -
sdw_vs   d10.61.65.0/24
          0.0.0.0/0      10.61.65.1      20

```

3. Make sure that the host can ping the Vserver interfaces and vice versa.

```

::> network ping -node <nodename> -destination <ip of host>

```

## New Vserver: SnapDrive User and Roles

SnapDrive must have rights on the Vserver to execute the various Data ONTAP API commands and to retrieve system information. SnapDrive requires two user accounts on the cluster:

- One with admin-level rights to the cluster Vserver that is hosting data (*vsadmin*)
- One with specific permissions to the admin Vserver for clusterwide operations, such as SnapMirror and licenses

1. Unlock vsadmin.
  - a. Set the password for vsadmin:

```

::> security login password -username vsadmin -vserver sdw_vs
Please enter a new password:
Please enter it again:

```

- b. Unlock the vsadmin account. It is locked by default:

```

::> security login unlock -username vsadmin -vserver sdw_vs

```

2. Create the clusterwide SnapDrive roles and user account.

The Vserver in this case is not the cluster Vserver that hosts data; instead, it is the admin Vserver that provides admin access to the cluster.

The roles required are:

```

::> security login role show -vserver <admin vservname> -role <snapdrive role name>

```

Example:

```

::> security login role show -vserver cluster1 -role sdw_role

```

Vserver	Role Name	Command/Directory	Access Query Level
cluster1	sdw_role	network	readonly
cluster1	sdw_role	snapmirror	readonly

cluster1	sdw_role	system health	readonly
cluster1	sdw_role	system license	readonly
cluster1	sdw_role	version	readonly

**Note:** SnapDrive might cache credentials; therefore, if an incorrect set of credentials is entered, adding correct credentials might fail. To correct this issue, restart the SnapDrive service to clear the credential cache.

- a. Create the roles with the following command (roles with spaces in the name should be enclosed in double quotation marks):

```
::>security login role create -role <snapdrive role name> -cmddirname [Command/Directory] -access [readonly|all]
```

**Example:**

```
::>security login role create -role sdw_role -cmddirname version -access readonly
```

- b. If a role already exists, use the Modify command to change the option. For example:

```
::>security admin role modify -role <snapdrive role name> -cmddirname network -access readonly -vserver <admin vservername>
```

- c. Create a new login account with ONTAPI® application permissions and password or domain authmethod and apply the newly created role to it:

```
::> security login create -username <username> -application ontapi -authmethod <password|domain> -role <role name> -vserver <admin vservername>
```

**Example:**

```
::> security login create -username snapdrive -application ontapi -authmethod password -role sdw_role -vserver cluster1
```

For more information on roles and security logins, refer to the System Administration Guide for your specific version of clustered Data ONTAP (available on the NetApp [Support](#) site).

## New Vserver: Create and Start SAN Services

For the host to communicate with the Vserver over a SAN protocol, the desired SAN protocol must be licensed, created, and started. Licenses are applied at the cluster level. SAN protocols are created and started at the Vserver level:

```
::>system license show
```

If the intended SAN protocol is not licensed, use the `license add` command.

1. Make sure that the desired SAN protocol is present:

```
::> iscsi show -vserver <cluster vservername>
::> fcp show -vserver <cluster vservername>
```

2. If the protocols are not present, create them:

```
::> iscsi create -vserver <cluster vservername> -status-admin up
::> fcp create -vserver <cluster vservername> -status-admin up
```

After the service is created, it needs to be started so that the cluster will listen for SAN traffic.

3. Start FCP or iSCSI if not already started:

```
::>vserver iscsi start
::>vserver fcp start
```

For more information about SAN protocols in clustered Data ONTAP, refer to the Fibre Channel and iSCSI Configuration Guide for your specific version of Data ONTAP (available at the NetApp [Support](#) site).

## Set Up with an Existing Vserver

Follow the steps, listed under section “New Vserver Setup” in the appendix, to create any missing components.

### Checklist for Verification

1. Verify what interfaces and volumes already exist. For SnapDrive to work with clustered Data ONTAP, you must have the following:
  - Vserver management LIF
  - SAN protocol LIF
  - Data volume

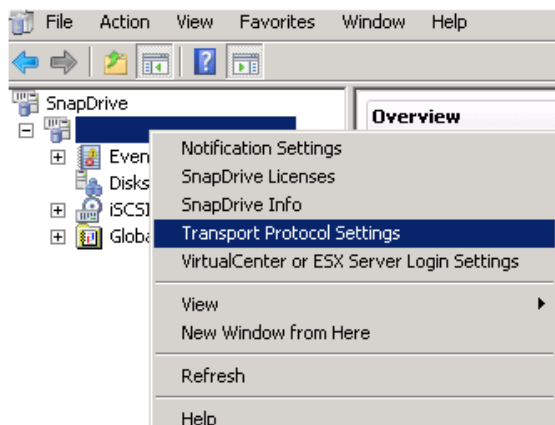
```
::> volume show -vserver <cluster vservername>  
::> net int show -vserver <cluster vservername>
```

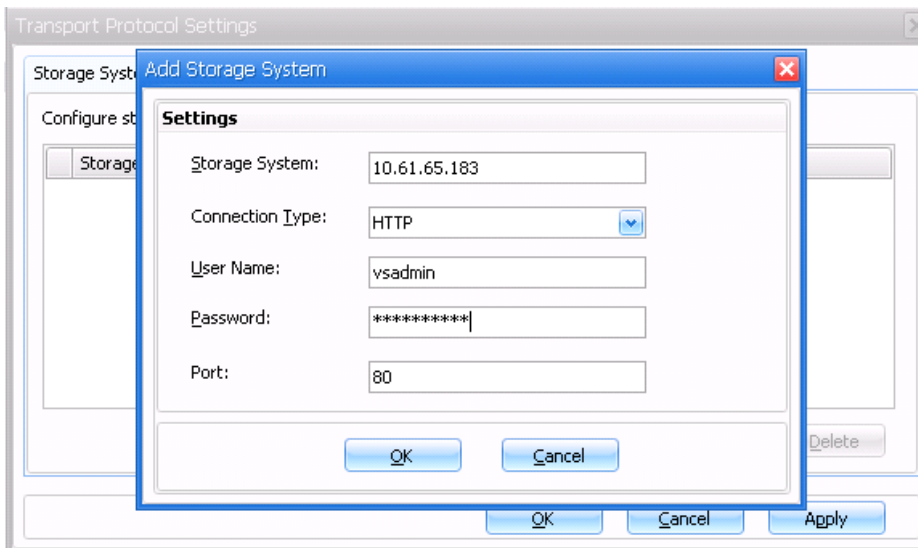
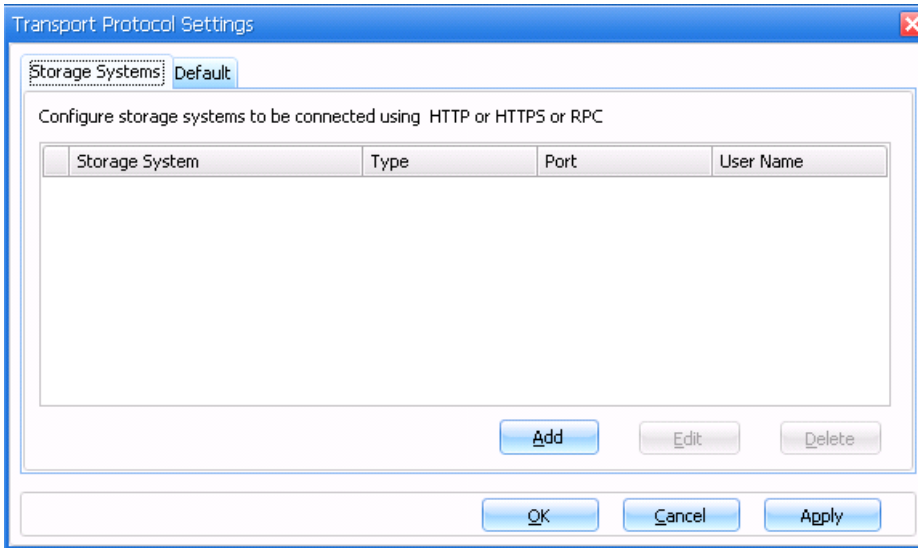
2. Create LIFs or volumes as needed.
3. Make sure that the host can ping the Vserver management LIF and that the cluster can ping the host.
4. Make sure that the SAN licenses are added.
5. Make sure that iSCSI and/or FCP services are created and enabled or started.
6. Make sure that the vsadmin account has been unlocked and that a password has been set.
7. Make sure that a default route exists for the Vserver.
8. Test IP connectivity from host and cluster.

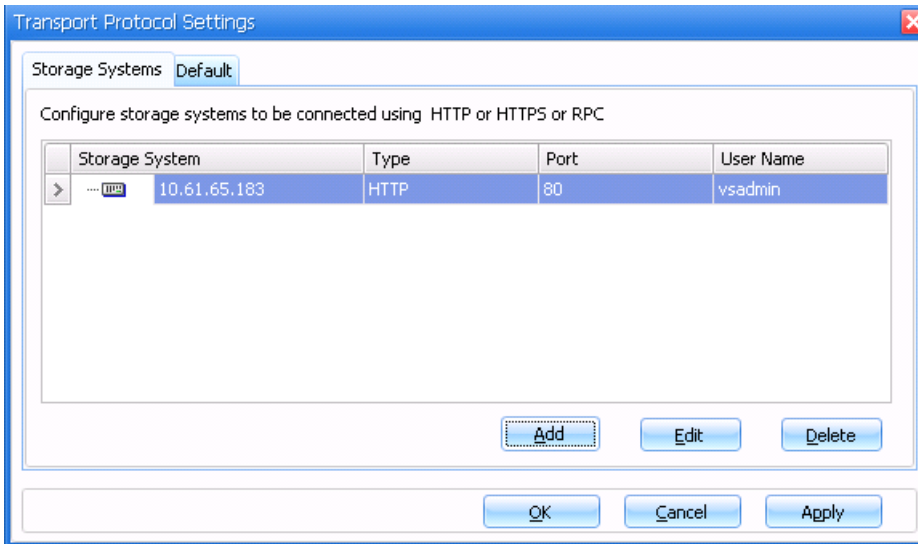
### Set Up SnapDrive from the Host Side

**Note:** RPC is no longer a supported protocol for SnapDrive when using clustered Data ONTAP. Only HTTP and HTTPS are supported with SnapDrive in clustered Data ONTAP.

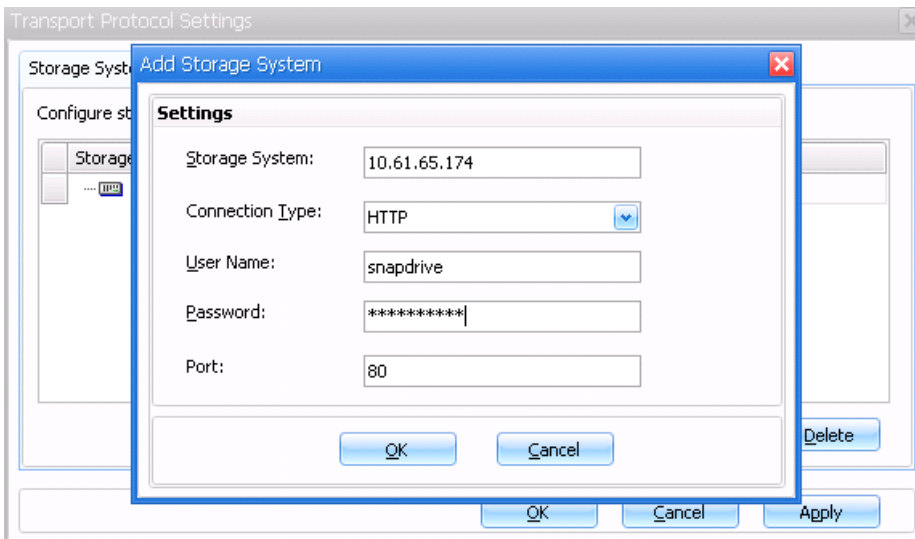
1. Install SnapDrive.
2. Make sure that iSCSI or FCP is configured or installed on the server or host.
3. Add the Vserver management IP to SnapDrive under Transport Protocol Settings by using HTTP or HTTPS and the vsadmin account.



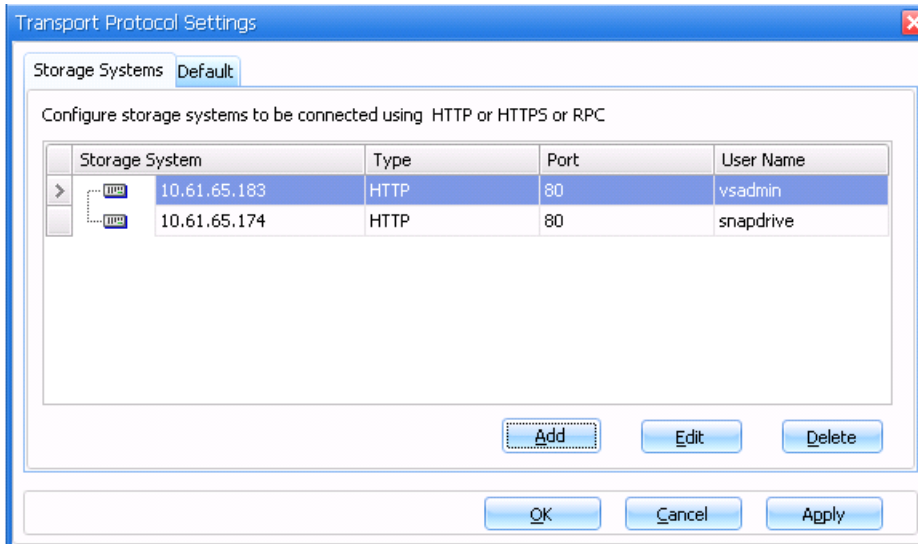




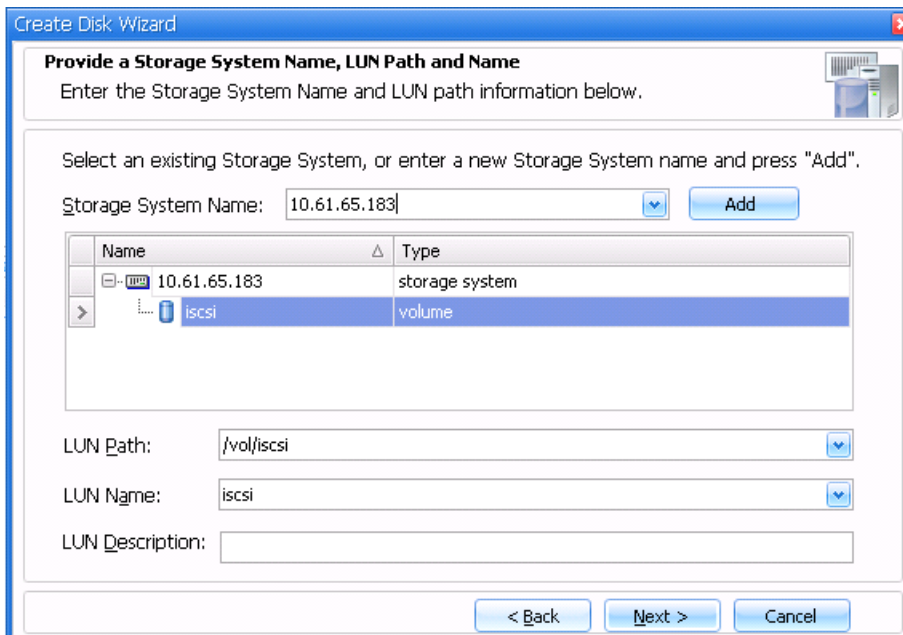
4. Add the cluster Vserver by using the created user or admin account. (If necessary, restart the SnapDrive service account to flush the credentials.)







5. Connect or create a disk by using the Vserver management IP.



## SnapDrive and Windows Server 2012 Behavior When a New Node Is Added to the Cluster

Consider three nodes named NODE1, NODE2, and NODE3 that are part of a cluster called W2012CLUS.

Also, assume that the nodes are configured as follows:

- NODE1 and NODE2 are part of cluster W2012CLUS
  - NODE1 is the owner of the following:
    - Cluster - W2012CLUS
    - "Available Storage" group
    - NODE1 and NODE2 have one dedicated disk each

- NODE1's dedicated disk has drive letter E:\
- NODE2's dedicated disk has drive letter F:\
- Now a new node, NODE3, is added to the cluster W2012CLUS using "Add Node Wizard." The wizard presents an option "Add all eligible storage to the cluster" to the user, which is selected by default.

## Behavior

- NODE3 is added to the cluster W2012CLUS

NODE1's dedicated disk (E:\) is converted as follows:

- Dedicated disk is converted to a clustered disk
- A cluster resource is created for the disk that is "Cluster Disk 1"
- Cluster resource is part of "Available Storage"
- Disk is online and accessible
- NODE2's dedicated disk (F:\) is converted as follows:
  - Dedicated disk is converted to a clustered disk
  - A cluster resource is created for the disk that is "Cluster Disk 2"
  - Cluster resource is part of "Available Storage" owned by NODE1
  - Disk is offline and not accessible by SnapDrive
- The dedicated disk of NODE2, that is, F:\, will neither be seen by NODE1 SnapDrive nor NODE2 SnapDrive
  - In NODE2, go to Disk Management and locate the dedicated disk. The disk will be offline. A mouse hover will give a tool tip "Reserved (The disk is offline because of policy set by an administrator)."
  - If you try to make the disk Online from the Disk Management console, you will see an error message: "The specified disk or volume is managed by the Microsoft Failover Clustering component. The disk must be in cluster maintenance mode and the cluster resource status must be online to perform this operation."
  - Also, SnapDrive gives out the following error for the disk in the debug logs: "The LUN may not have a file system created on it, or it may not be formatted with the NTFS file system. If it is a cluster system this situation can be caused by the LUN being offline, or not owned by this computer."

## Resolution

To bring the "Cluster Disk 2" on line, perform the following steps:

1. Add the i-SCSI Initiator of NODE1 to the LUN using NetApp FilerView\Cluster Element Manager
2. Bring the associated Cluster Resource "Cluster Disk 2" on line in Failover Cluster Manager.
3. Refresh "Disks" in SnapDrive MMC Console and you will be able to see the "Cluster Disk 2" in SnapDrive.
4. Repeat the same operation for "Cluster Disk 1" to see the disk in SnapDrive.

## Acknowledgement

Justin Parisi, Escalation Engineer, NetApp

## Version History

Version	Date	Document Version History
Version 1.0	February 2012	Initial release.
Version 1.0.1	January 2013	Updated with section 21 for Windows Server 2012 support.

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

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

Go further, faster®

© 2013 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, ASUP, AutoSupport, DataFabric, Data ONTAP, FlexClone, ONTAPI, SnapDrive, SnapManager, SnapMirror, SnapRestore, Snapshot, and SnapVault are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Microsoft, SharePoint, SQL Server, Windows, Windows Server, and Windows Vista are registered trademarks and Hyper-V is a trademark of Microsoft Corporation. ESX, vMotion, and VMware are registered trademarks and ESXi is a trademark of VMware, Inc. Catalyst and Cisco Nexus are registered trademarks of Cisco Systems, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-4000-0113