



Technical Report

SnapDrive for Windows for Clustered Data ONTAP Used in SAN Environments

Best Practices Guide

Microsoft Business Unit, NetApp
December 2014 | TR-4228

Abstract

This technical report provides best practices guidelines for implementing NetApp® SnapDrive® 7.x for Windows® in NetApp clustered Data ONTAP® 8.2 and later used in SAN environments.

TABLE OF CONTENTS

1	Overview	4
2	Purpose and Scope	4
3	SnapDrive for Windows Key Features	4
4	Clustered Data ONTAP	5
4.1	Characteristics of Clustered Data ONTAP	5
4.2	What's New in Clustered Data ONTAP 8.3	6
5	Key Features in Clustered Data ONTAP	6
5.1	Offload Data Transfer	6
5.2	Storage Quality of Service	7
6	What's New in SnapDrive 7.1 for Windows	8
6.1	Clone of Clone	8
6.2	Split Clone	9
6.3	Dedicated Disk Optimization	10
6.4	Thin Provisioning for LUNs	10
6.5	MetroCluster Support	10
7	SnapDrive for Windows Architecture	10
8	Licensing Requirements for SnapDrive 7.x for Windows	11
9	Deploy the Storage System for SAN Environments	11
10	Protocols Supported in SnapDrive for Windows	12
11	Upgrading SnapDrive for Windows	12
12	Data ONTAP DSM	12
13	Asymmetric Logical Unit Access	13
14	LUN Management in SnapDrive for Windows	13
14.1	LUN Management During Volume Moves	14
14.2	Space Reclamation	15
15	Snapshot Management	15
15.1	Snapshot Scheduling	15
15.2	Storage Failover	16
16	DataFabric Manager and Role-Based Access Control Support	16
17	AutoSupport	16

18 Data Protection	16
18.1 SnapMirror	16
18.2 SnapVault	17
18.3 Restore	18
18.4 IPv6.....	18
19 SnapManager Dependency on SnapDrive	18
20 Windows Server 2012 and Windows Server 2012 R2 Support	19
20.1 Introduction	19
20.2 Prerequisites	19
20.3 Features Overview	20
20.4 Windows Server 2012 and Windows Server 2012 R2 Feature Not Supported in SnapDrive.....	21
20.5 Windows Server 2012 Virtual Machine Support for ESXi Environments	22
21 VMware Support	22
22 Performance Tuning.....	23
Appendix.....	23
SnapDrive and Windows Server 2012 Behavior When a New Node Is Added to the Cluster	23
Version History	24

LIST OF TABLES

Table 1) Redundancy achieved in each layer.....	5
Table 2) Supported protocols.	12

LIST OF FIGURES

Figure 1) Offloaded data transfer.	7
Figure 2) SVM with its own QoS policy.....	7
Figure 3) Clone of clone and split clone.	8
Figure 4) SnapDrive for Windows architecture.	11
Figure 5) Example of ALUA optimized path.....	13
Figure 6) Transport protocol settings page showing Cserver and SVM details.	14
Figure 7) Storage failover.	16
Figure 8) SnapDrive and SnapVault integration.	17
Figure 9) SQL Server virtual machine hosted on clustered Data ONTAP.	19
Figure 10) SQL Server database hosted on VMDKs by using SnapDrive and NetApp VSC.....	22

1 Overview

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

SnapDrive manages LUNs on a storage system, making those 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).

SnapDrive can also provision SMB 3.0 shares to support Microsoft Hyper-V® over SMB and Microsoft SQL Server® over SMB workloads.

2 Purpose and Scope

SnapDrive for Windows performs storage management tasks, enables application-consistent backups, and restores data by integrating with NetApp SnapManager® products. It also enables replication of NetApp Snapshot® copies to remote storage for both SAN and SMB environments.

The best practices guidelines in this document apply to SnapDrive in NetApp clustered Data ONTAP 8.2 and later used in SAN environments.

For best practices guidelines in SMB environments, refer to [TR-4218: SnapDrive for Windows Best Practices Guide for Clustered Data ONTAP for NAS](#).

3 SnapDrive for Windows Key Features

SnapDrive for Windows supports the following features in clustered Data ONTAP environments:

- Enables online storage configuration, LUN expansion, and streamlined management
- Integrates Snapshot technology with Microsoft Volume Shadow Copy Service (VSS) framework, which enables the creation of point-in-time images of data stored on LUNs
- Integrates with Microsoft remote VSS framework to perform backup and restore of SMB 3.0 shares hosted on NetApp storage systems that run clustered Data ONTAP 8.2 and later
- Enables support for Microsoft cluster configurations
- Supports iSCSI, Fibre Channel (FC), and Fibre Channel over Ethernet (FCoE) on Windows and VMware® environments
- Supports virtual Fibre Channel (vFC) adapters for guest virtual machines (VMs) in Windows Server 2012 and Windows Server 2012 R2 environments
- Supports raw device mapping (RDM), Virtual Machine Disk (VMDK) over NFS, and Virtual Machine File System (VMFS) datastores in VMware environments
- Enables NetApp SnapVault® integration natively
- Supports Group Managed Service Accounts in Windows Server 2012
- Supports IPv6
- Supports Data ONTAP PowerShell cmdlets to execute SAN and SMB workflows

4 Clustered Data ONTAP

Storage controllers running clustered Data ONTAP are referred to as *nodes*. These nodes are joined in a clustered system. The nodes in the cluster continuously communicate with each other, coordinate cluster activities, and transparently move data between nodes.

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

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

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

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

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

4.1 Characteristics of Clustered Data ONTAP

Transparent Data Mobility

With clustered Data ONTAP, you can move volumes from one aggregate to another within an SVM. 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.

Achieving High Availability

To create a highly available, redundant environment, make sure that redundancy is addressed in the layers described in Table 1.

Table 1) Redundancy achieved in each layer.

Layer	Description
Sessions from the host to the storage subsystem	Windows multipath input/output (MPIO), Data ONTAP device-specific module (DSM)
Server	Windows failover cluster manager
Ethernet	Two Ethernet switches
Fibre Channel	Two fabric switches
Storage systems	Two-node clustered Data ONTAP system (minimum)
LIFs in the storage systems	Multiple LIFs mapped to the same port

For more information, refer to [TR-4080: Best Practices for Scalable SAN in Clustered Data ONTAP 8.2](#).

4.2 What's New in Clustered Data ONTAP 8.3

MetroCluster Support

NetApp MetroCluster™ support for clustered Data ONTAP 8.3 provides a highly efficient and reliable disaster recovery solution. It works across applications; in virtual environments; and across storage types, protocols, and tiers. MetroCluster is the solution that provides continuous availability for critical applications for which you can never afford planned or unplanned downtime.

The minimum configuration for MetroCluster is a disaster recovery group that consists of one HA pair at each site, for a total of four nodes (controllers). Each cluster is an active-active HA pair, so all nodes serve clients at all times. This solution can stretch across citywide or metrowide deployments, up to a maximum distance of 124 miles (200km). This capability enables a level of availability that goes beyond the HA features of a local cluster, which makes MetroCluster a highly versatile solution.

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

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

5 Key Features in Clustered Data ONTAP

This section describes the key features of Data ONTAP 8.2 and later.

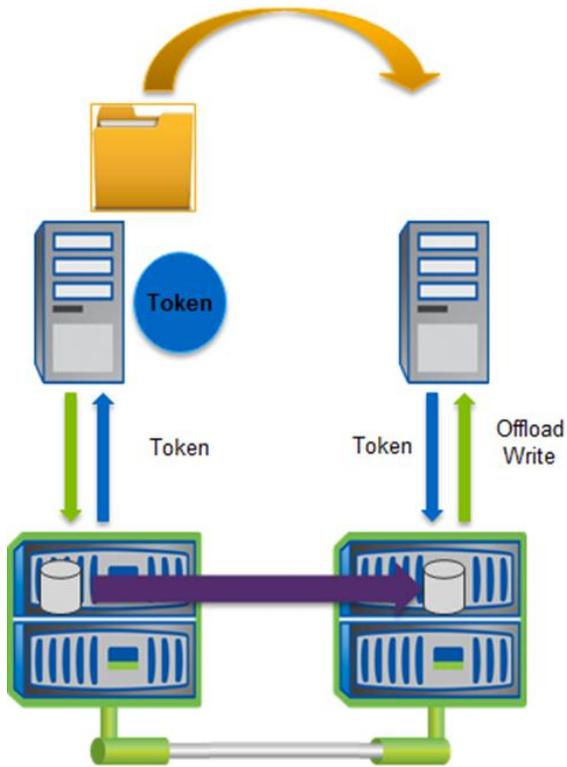
5.1 Offload Data Transfer

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

With Windows Server 2012 and Windows Server 2012 R2, Microsoft introduced a copy offload mechanism, which allows you to offload the activity of copying files between two servers to the storage system.

Before proceeding with the copy operation on the host, make sure that copy offload settings are configured on the storage system. SnapDrive 7.x for Windows does not enable copy offload settings on the storage system directly. After LUNs and shares are provisioned by using SnapDrive, activities such as storage live migration and file copy are offloaded to the storage system. SnapDrive does not participate in these operations.

Figure 1) Offloaded data transfer.



5.2 Storage Quality of Service

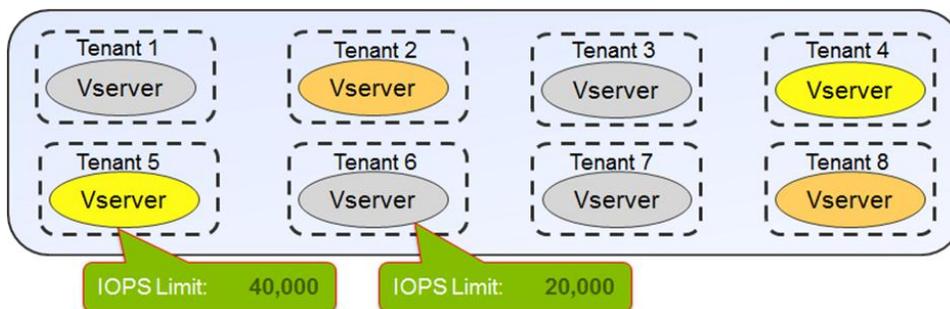
Storage quality of service (QoS) in clustered Data ONTAP 8.2 and later provides the ability to group storage objects and set throughput limits on the group. With this ability, you can separate workloads by organization, application, business unit, or production or development environments.

In enterprise environments, storage QoS helps to achieve the following:

- Prevents user workloads from affecting each other
- Protects critical applications that have specific response times that must be met in IT-as-a-service environments
- Prevents tenants from affecting each other
- Avoids performance degradation with the addition of each new tenant

QoS allows you to limit the amount of I/O sent to an SVM, a flexible volume, a LUN, or a file. I/O can be limited by the number of operations or the raw throughput.

Figure 2) SVM with its own QoS policy.



After you configure SVMs or volumes for QoS, add details of those SVMs on the host. You can then use SnapDrive 7.x for Windows to provision storage from those SVMs onto the host system or the guest virtual machines. For Microsoft application-specific QoS best practices, refer to the respective SnapManager best practices guides.

6 What's New in SnapDrive 7.1 for Windows

SnapDrive 7.1 for Windows introduces the following new features.

6.1 Clone of Clone

The clone-of-clone feature enables you to provision a volume, create a Snapshot copy of that volume, and then clone the volume by using the Snapshot copy. Data ONTAP supports the clone-of-clone feature. Clustered Data ONTAP supports two clones: SIS clone and volume clone (NetApp FlexClone® technology).

Data ONTAP creates a new volume that hosts the data from the Snapshot copy from which it was created. This volume clone is tightly bound to the parent volume because the underlying blocks are shared. On top of this volume clone, you can create another Snapshot copy and mount that Snapshot copy as a volume clone. SnapManager for SQL Server uses the clone-of-clone functionality to clone the data for testing and development purposes.

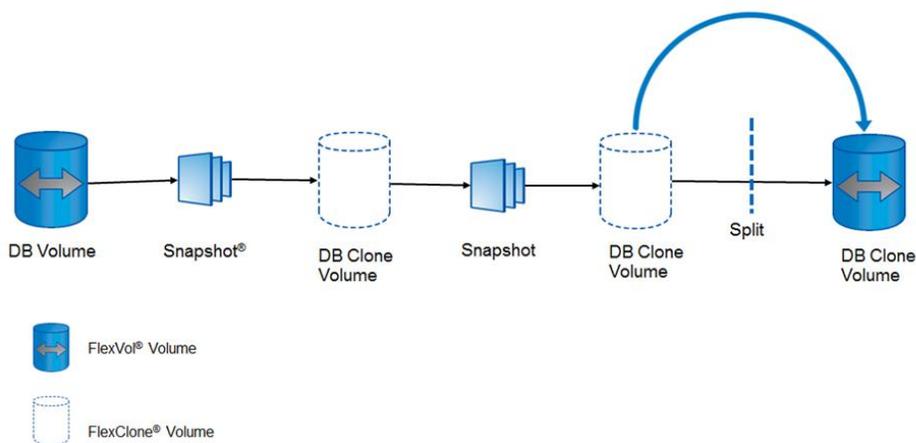
Note: Creating a LUN clone from SnapDrive is not supported.

When you use the clone-of-clone feature, consider the following information:

- By default, SnapDrive supports only two levels of clones; however, in the case of VMDK on VMFS, more than two levels can be created.
- Make sure that there is adequate space in the aggregate for clone creation.
- SnapDrive doesn't allow LUN creation on a clone volume.
- The clone-of-clone operation cannot be performed in a pass-through disk connected to an HA VM.

Figure 3 shows the details of the clone-of-clone and split-clone features.

Figure 3) Clone of clone and split clone.



SnapDrive 7.1 for Windows also provides new sdcli commands that support the clone-of-clone feature. Use the following commands to initiate the clone-of-clone operation:

```
sdcli snap mount -r host3 -k j -s Jun_14_09 -d t
```

```
sdcli disk connect -d s -dtype shared -p storage2:/vol/sd_vds_only/mktng.lun -I host3
10:00:00:00:C9:2B:FD:1B host3 10:00:00:00:C9:2B:FD:1C host4 10:00:00:00:C9:2B:FD:12 host4
10:00:00:00:C9:2B:FD:11 -e "tech_mktng" -c "mktng"
```

6.2 Split Clone

The split-clone feature enables you to split the volume clone from the parent volume and make the volume clone an independent NetApp FlexVol® volume. All SnapManager products make use of the split-clone feature. When the Snapshot copy is mounted as a FlexClone volume, the volume clone depends heavily on the parent volume because the blocks are shared. The split-clone feature allows you to remove the dependency of the FlexClone volume by splitting the clone from the parent. Data ONTAP supports the split-clone feature.

When you use the split-clone feature, consider the following information:

- During a volume clone split operation, LUN provisioning and Snapshot management operations are not supported on the effected volume.
- Run the `sdcli clonesplit status` command to determine the status of the split-clone operation.
- Make sure that the aggregate has adequate space for the split-clone operation.
- Because split-clone operations take considerable time (a 1GB volume clone split operation takes 5 minutes, for example), make sure that you account for the required time.

SnapDrive 7.1 for Windows also provides new `sdcli` cmdlets to support the split-clone feature. Use the following command to estimate the space in the aggregate for the volume clone split operation:

```
sdcli clone_split clone_estimate -D F:

Volume Name           = clone1
Mount Point           = F
Space required in Aggregate = 1.25GB
Space in Aggregate    = 714.62GB
Can Split              = Yes
Operation Status      = Complete
The operation completed successfully.
```

The following command starts the volume clone split for the FlexClone volume:

```
sdcli clone_split clone_start -D F:
WARNING: Performing Volume Split on the Clone will delete all the snapshots taken on the volume clone.
Do you want to Start the Split Clone operation?[Y/N] : Y

Volume Name           = clone1
Mount Point           = F
Operation Status      = Complete
The operation completed successfully.
```

The following command displays the status of the volume clone split operation:

```
sdcli clone_split clone_status -D F:

Volume Name           = clone1
Mount Point           = F
Operation Status      = In Progress
The operation completed successfully.
```

6.3 Dedicated Disk Optimization

With SnapDrive 7.1 for Windows, you install SnapDrive only on the cluster node on which the dedicated disk is created. This is a change from previous releases in which SnapDrive for Windows was installed on every node in the cluster.

Note: This change is applicable to Windows Server failover clustering environments, and it is specific to dedicated disks.

For this optimization to work correctly, you must configure the transport protocol settings on the cluster node on which the dedicated disk is created.

This enhancement applies to Microsoft Exchange database availability groups, SQL Server availability groups, and mixed storage environments in which SnapDrive and SnapManager will be used only in certain nodes in the cluster.

6.4 Thin Provisioning for LUNs

In earlier SnapDrive for Windows releases, thin provisioning for LUNs was possible only when the space guarantee property was set to none. Starting in SnapDrive 7.1 for Windows, you can initiate thin provisioning regardless of the space guarantee property value.

With SnapDrive 7.1, a thin-provisioned LUN can be created by using the new `-thin` parameter in the `sdcli create disk` command. A thin-provisioned LUN can also be created through the Create Disk wizard.

6.5 MetroCluster Support

MetroCluster in clustered Data ONTAP is a scalable solution that provides cluster-local HA and cross-site synchronous disaster recovery. SDW works seamlessly after the new cluster begins functioning. After the switchover operation completes, SDW enumerates the disks with new SVMs, and all the SnapDrive related operations are performed.

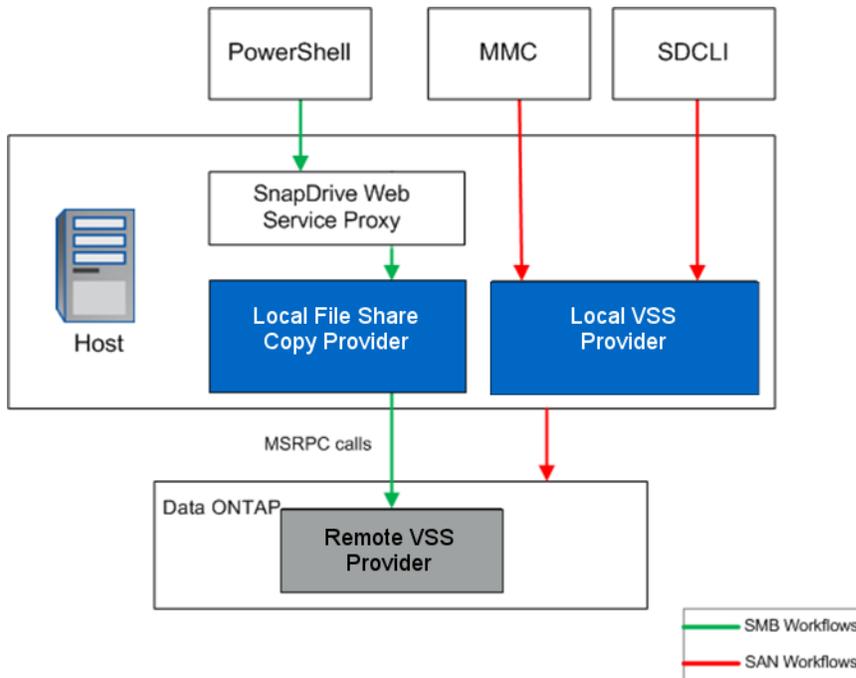
Note: When you use VMDK disks, you must restart SDW and the Virtual Storage Console (VSC) service.

7 SnapDrive for Windows Architecture

SnapDrive 7.x for Windows supports both SAN and SMB workflows. All SAN-based workflows can be performed by using the SnapDrive GUI or `sdcli` commands. With SMB workflows, the web service proxy passes requests such as backup, restore, and so on, to the local file share copy provider, which communicates to the remote VSS provider in the Data ONTAP system by using MSRPC calls. If a SAN operation is initiated, the call is directly passed to the VSS framework and reaches the VSS hardware provider that resides on the host. This in turn communicates to Data ONTAP for SAN-related operations.

Any SMB or SAN backup, restore, or replication operations initiated by using SnapManager products, such as SnapManager for Hyper-V or SnapManager for SQL Server, reach the SnapDrive web service proxy layer directly and are redirected accordingly.

Figure 4) SnapDrive for Windows architecture.



8 Licensing Requirements for SnapDrive 7.x for Windows

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

- FlexClone (for mount operations)
- CIFS (Hyper-V over SMB or SQL Server over SMB workloads)
- iSCSI
- FCP
- NetApp SnapRestore®
- NetApp SnapMirror® (optional)
- SnapVault (optional)
- SnapManager suite (if you are using storage system–based licensing)
- Data ONTAP device-specific module (DSM) (for multipathing)

Note: All the SnapManager products for Windows environments—SnapManager for Hyper-V, SnapManager for SQL Server, SnapManager for SharePoint®, and SnapManager for Exchange—communicate directly with SnapDrive for Windows for the host license check.

9 Deploy the Storage System for SAN Environments

For information about how to deploy the storage system for a SAN environment, refer to the following resources:

- [SnapDrive 7.x for Windows Installation Guide](#) for information about how to install and upgrade to SnapDrive 7.1 for Windows
- [SnapDrive 7.x for Windows Quick Start Guide for Clustered Data ONTAP](#) for information about how to deploy SnapDrive for Windows in a clustered Data ONTAP environment

10 Protocols Supported in SnapDrive for Windows

Table 2 lists the protocols supported in SnapDrive for Windows.

Table 2) Supported protocols.

Parameter	Data ONTAP 7-Mode Storage Systems	Clustered Data ONTAP Storage Systems
Storage protocol	<ul style="list-style-type: none">SAN over iSCSI or FC or FCoEFor ESX environments: VMDK over NFS and VMDK over VMFS	<ul style="list-style-type: none">SAN over iSCSI or FC or FCoEFor ESX environments: VMDK over NFS and VMDK over VMFSSMB 3.0 or Windows Server 2012 and Windows Server 2012 R2 and clustered Data ONTAP 8.2 and later
Network protocol	IPv4 and IPv6	IPv4 and IPv6
Transport protocol	RPC, HTTP, and HTTPS	HTTP and HTTPS

11 Upgrading SnapDrive for Windows

If you're adding a clustered Data ONTAP storage system to the same host that is using Data ONTAP operating in 7-Mode systems, to enable multipathing, before you upgrade SnapDrive, you must upgrade the DSM on the host system to allocate storage from clustered Data ONTAP storage systems.

12 Data ONTAP DSM

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 x64 and Windows Server 2012 and Windows Server 2012 R2 servers include support for Microsoft MPIO.

Data ONTAP DSMs for Windows MPIO help NetApp storage systems integrate with Microsoft MPIO on Windows servers and provide high availability to applications by using path-failover methods. Data ONTAP DSMs determine all of the paths that point to the same LUN so that MPIO can group them into the virtual disk that the Windows 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. Data ONTAP uses Asymmetric Logical Unit Access (ALUA) to identify optimized paths. There can be multiple active paths and multiple passive paths. You can have a maximum of 32 paths to a LUN. This maximum applies to any mix of FC and iSCSI paths. If all the active paths fail, the DSM automatically switches to the passive paths, maintaining the host's access to its storage.

The DSM 4.1 and DSM 4.1P1 installer includes all settings and files required for the driver. This eliminates the requirement for Windows Host Utilities with Data ONTAP DSM 4.1 and DSM 4.1P1. The installer includes the MBRAAlign tool and the `LinuxGuestConfig.iso` in the installation package. The installer displays a dialog box that prompts you to install the Hyper-V Guest Utilities; the default value is No. Selecting Yes installs the MBRAAlign tool and `LinuxGuestConfig.iso` in the `INSTALLDIR` directory (the default is `%systemdrive% Program Files\NetApp\MPIO`).

For more information about installation and configuration of Data ONTAP DSM 4.1 and DSM 4.1P1, refer to [Data ONTAP DSM 4.1 for Windows 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.

For more details about installation steps, refer to the [SnapDrive for Windows Installation Guide](#) and the [NetApp Interoperability Matrix Tool](#) (IMT).

13 Asymmetric Logical Unit Access

ALUA support is available on all clustered Data ONTAP configurations with Data ONTAP DSM as well as with the Microsoft DSM.

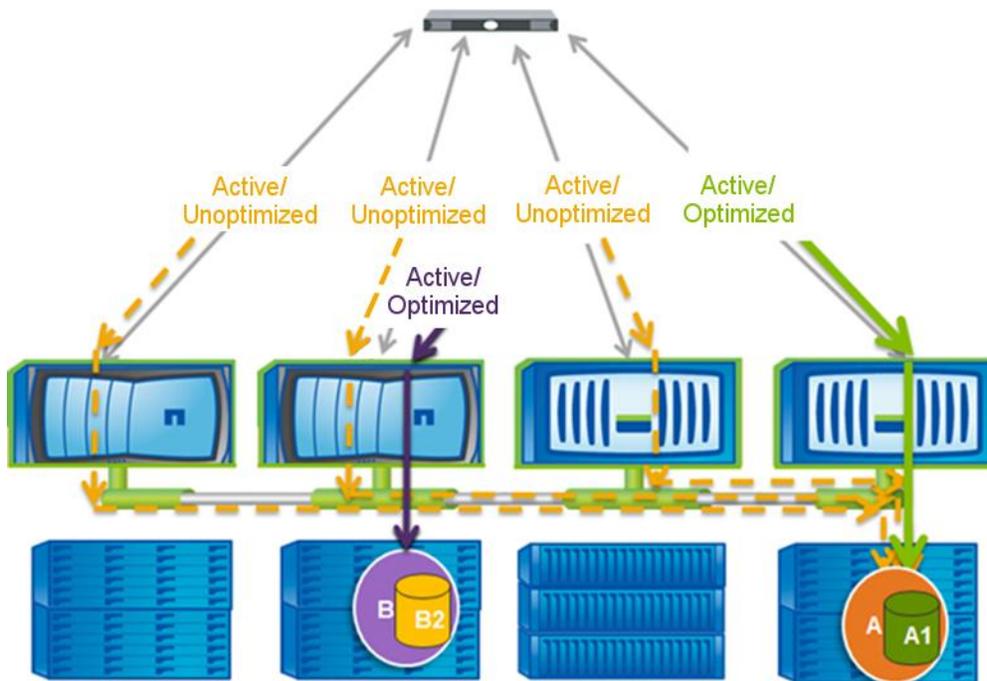
Hosts use ALUA to determine the state of a specified path. The state of a specified path can be one of the following:

- **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.

By default, clustered Data ONTAP enables ALUA. When you create a shared disk, take into account the following consideration with respect to the host server OS that is used:

- Windows 2008 and later should be used for clustered Data ONTAP shared LUNs.

Figure 5) Example of ALUA optimized path.



14 LUN Management in SnapDrive for Windows

In clustered Data ONTAP, you can isolate a single workload within an SVM. An SVM functions as a logical container that has its own namespace. After you create the SVM, 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.

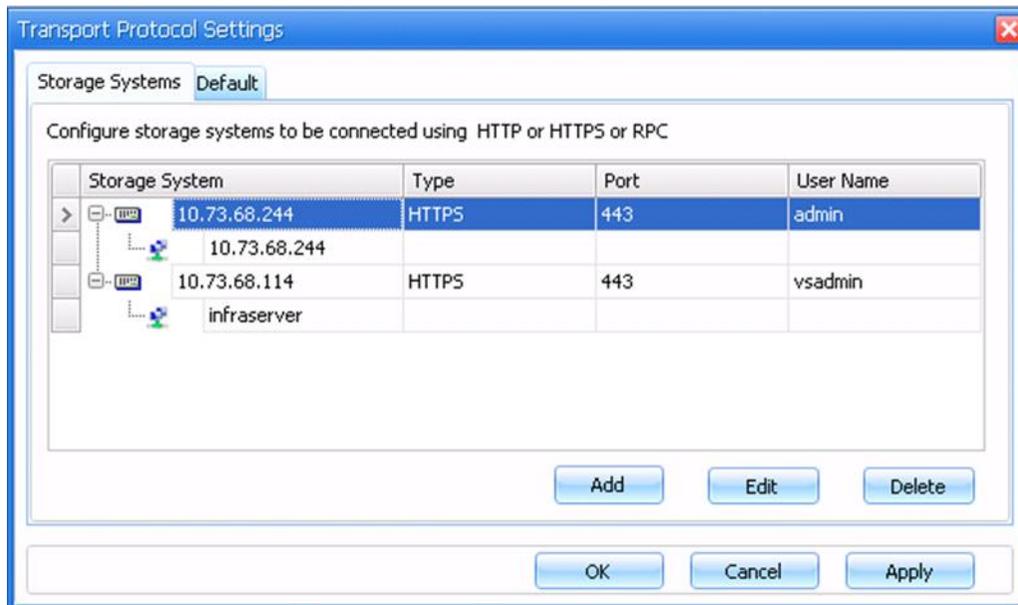
Before performing any LUN management task, you must add the management LIF that was created during the SVM creation. This LIF communicates internally with other data LIFs such as iSCSI or FC LIFs; if the host system is hosting SMB workloads in addition to SAN, it is required to segregate SAN and SMB traffic by using two different LIFs.

Best Practice

At least one data LIF and one management LIF per SVM are required for SnapDrive to perform LUN management tasks. Make sure that the aggregate that is hosting the LUN is part of the aggregate list of the SVM.

Before performing any operation, the SVM management LIF located in SnapDrive > Transport Protocols Settings > Storage Systems must be entered. Also, with SnapDrive 7.x for Windows, it is no longer a requirement to enter the cluster credentials for replication by using SnapMirror in Data ONTAP 8.2 and later.

Figure 6) Transport protocol settings page showing Cserver and SVM details.



14.1 LUN Management During Volume Moves

Clustered Data ONTAP supports volume moves from one aggregate to another aggregate anywhere within the cluster. Volume moves are performed when cluster expansion is needed, 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 from NetApp OnCommand® System Manager). Volume moves are nondisruptive and provide uninterrupted access to the hosts that are connected to the storage by using any protocol, such as iSCSI, NFS, CIFS, FC, or FCoE, during the volume move.

This means that you 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.

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.

14.2 Space Reclamation

Space reclamation is a feature in SnapDrive that is used to communicate with the NetApp storage system to free up blocks in a LUN that are marked as “free” by the NTFS metadata. Running space reclamation periodically is necessary to reclaim the free space in the LUN. Use the following command to reclaim free space:

```
SDCLI spacereclaimer start [-m <MachineName>] -d <Mount Point> [-t <TimetoRun>]
```

Before you run space reclamation on Cluster Shared Volume (CSV) LUNs in Windows Server 2012, make sure that CSV LUNs are in redirected mode.

When you use the space-reclamation feature, consider the following information:

- Space reclamation is not supported on disks created outside of SnapDrive.
- Space reclamation is not supported on VMDK files created on VMFS datastores.

15 Snapshot Management

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

When SnapDrive for Windows (SDW) initiates a Snapshot copy on a LUN present in a Data ONTAP 7-Mode system, a LUN clone operation is initiated. This process uses FlexClone technology in the back end. In addition, to restore 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: Operations such as connecting to a disk within a volume or to a disk within a Snapshot copy require a FlexClone license.

15.1 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.

With clustered Data ONTAP systems, make sure that the Snapshot scheduling is turned off at the volume level before you perform storage management tasks on the host.

Best Practice

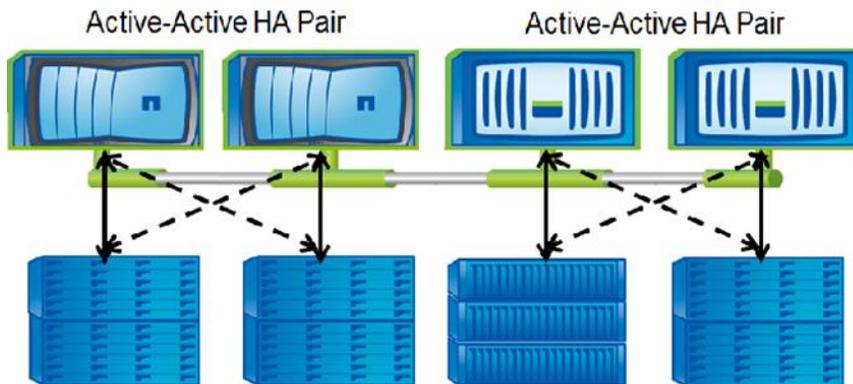
Because the SnapManager suite of products (SnapManager for SQL Server, SnapManager for SharePoint, SnapManager for Exchange, and SnapManager for Hyper-V) uses SnapDrive for application-consistent Snapshot copies, NetApp recommends making sure that there are minimal overlaps when these application-specific Snapshot copies are initiated through their respective products. This reduces the performance overhead on the cluster storage systems.

15.2 Storage Failover

To accommodate storage failover, two storage controllers (nodes) in the same cluster are connected together as a storage failover (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 7) Storage failover.



16 DataFabric Manager and Role-Based Access Control Support

SnapDrive for Windows does not support role-based access control (RBAC) for clustered Data ONTAP. However, a workaround is available by using the RBAC User Creator tool, which is available on the [NetApp Support](#) site.

17 AutoSupport

SnapDrive supports NetApp AutoSupport™ (ASUP™) diagnostics for systems based on both Data ONTAP operating in 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 SnapManager specific ASUP data such as SnapManager for SQL Server events, you must have ASUP enabled.

18 Data Protection

18.1 SnapMirror

Clustered Data ONTAP 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 SVMs in the same cluster

Note: You cannot establish a SnapMirror relationship between a Data ONTAP 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.

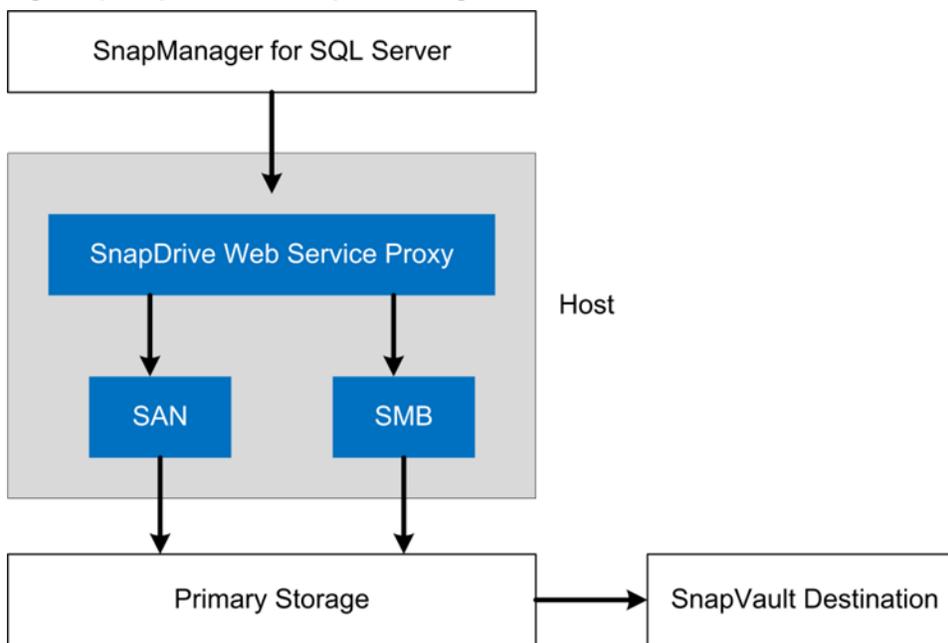
In SnapDrive 6.5, it was mandatory to provide cluster credentials to initiate SnapMirror updates when connecting to clustered Data ONTAP 8.1.x systems. However, SnapDrive 7.x for Windows does not require cluster VSS credentials to be provided to allow SnapMirror updates when connecting to clustered Data ONTAP 8.2 or later systems. The earlier version of clustered Data ONTAP still requires cluster credentials, however.

18.2 SnapVault

SnapDrive 7.x for Windows introduces native SnapVault integration. This means that in clustered Data ONTAP 8.2 and later environments, NetApp OnCommand Unified Manager is no longer required to configure and update SnapVault datasets. SnapVault can be configured by using OnCommand System Manager or Windows PowerShell cmdlets from a Windows host.

Note: While configuring SnapVault, make sure that the SnapVault destination details are added in the transport protocol settings in SnapDrive.

Figure 8) SnapDrive and SnapVault integration.



For SnapVault deployment instructions and best practices in Microsoft application environments, refer to the documentation available on the [Field Portal](#).

18.3 Restore

In scenarios in which data retrieval is required from an archive destination, SnapManager can request a SnapVault restore from SnapDrive. SnapDrive initiates a LUN- or file-level restore of the database from a Snapshot copy present in the SnapVault destination.

Note: LUNs created on a SnapVault source volume are displayed in SDW as dedicated LUNs.

Note: SnapVault can be executed at the volume level only. Qtree-level SnapVault operations are not supported.

Note: A SnapRestore license is required on the SnapVault destination for restoring Snapshot copies.

18.4 IPv6

Clustered Data ONTAP 8.2 and later supports IPv6. SnapDrive 7.x for Windows supports IPv6 in all its workflows that require adding an IP address or host name, such as adding a storage system in the transport protocol settings. Use the following command to enable IPv6 in clustered Data ONTAP 8.2 and later:

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

Note: Make sure that the DNS server has either the IPv6 or the IPv4 address of the host system, but not both.

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

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

Note: NetApp VSC 4.2.x and 5.0 do not support IPv6.

For more information about IPv6, refer to the [Clustered Data ONTAP Network Management Guide](#).

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 to create application-consistent Snapshot copies.

Best Practice

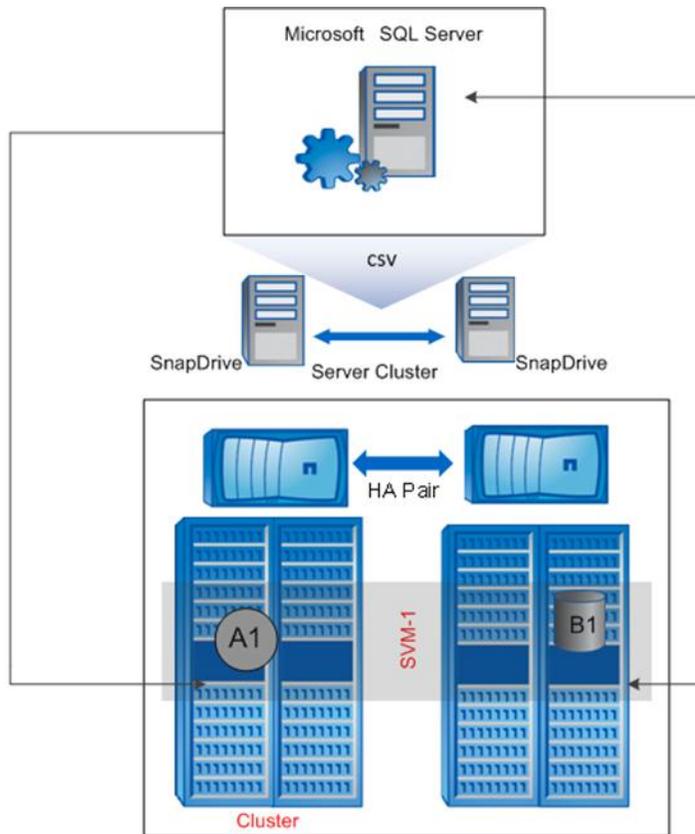
Because the SnapManager suite of products uses SnapDrive for application-consistent Snapshot copies, NetApp recommends making sure that no overlaps occur when these application-specific Snapshot copies are initiated by using their respective products. This reduces the performance overhead on the storage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

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

Best Practice

Store the applications and the VMs that host them in a single SVM. This facilitates secure isolation and simplifies backups and restores. The volumes in this SVM can have a SnapMirror relationship with another SVM containing an identical set of volumes. This also simplifies the disaster recovery strategy.

Figure 9) SQL Server virtual machine hosted on clustered Data ONTAP.



	Volume in SVM-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 SVM-1.
SVM-1	SVM-1 hosts all the data volumes that are specific to SQL Server (VMs and database).
SnapDrive	SnapDrive for Windows installed on the host server.
	SnapManager for SQL Server and SnapDrive for Windows installed on the guest VM.

20 Windows Server 2012 and Windows Server 2012 R2 Support

20.1 Introduction

SnapDrive for Windows supports Windows Server 2012 and Windows Server 2012 R2.

Note: SnapDrive 7.x for Windows and the supported SnapManager products that are aligned with SnapDrive 7.x are mandatory when you deploy a Windows Server 2012 and Windows Server 2012 R2 environment with Data ONTAP 8.2 and later.

20.2 Prerequisites

The following prerequisites are required to support Windows Server 2012 and Windows Server 2012 R2:

- **MSDSM or Data ONTAP DSM 4.1 and 4.1P1.** You can choose to install either MSDSM or Data ONTAP DSM 4.1 and 4.1P1 to enable multipathing.
- **Windows Host Utilities Kit 6.0.2.** This is required to disable native space reclamation carried out by Windows Server 2012 and 2012 R2.

Note: Windows Host Utilities Kit is not required if you have already installed Data ONTAP DSM 4.1 for Data ONTAP 8.2 and later environments.
- **.NET Framework 4.0 and .NET Framework 3.5 SP1.** SnapDrive for Windows requires .NET Framework 4.0 and .NET Framework 3.5 SP1 for installation.

20.3 Features Overview

SnapDrive 7.x for Windows supports all major SAN and SMB-based features in Windows Server 2012 and Windows Server 2012 R2. This section provides an overview of all the features and best practices for SAN environments.

CSVFS Support

The Cluster Shared Volumes proxy file system (CSVFS) is supported in SnapDrive for Windows 6.5 and later. In Windows Server 2012 and Windows Server 2012 R2, CSVs 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 that functions as a layer of abstraction above the NTFS for the storage volume. As a result, simultaneous reads and writes can be performed on the CSV LUN from different nodes. For more information about CSVFS, refer to the Microsoft TechNet article [Use Cluster Shared Volumes in a Failover Cluster](#).

A CSVFS volume has two volume GUIDs, including:

- **NTFS volume GUID.** The NTFS volume GUID is assigned when a disk is created and partitioned with NTFS before it is added to the CSV.
- **CSV volume GUID.** The CSV volume GUID is assigned when a disk is added to the CSV.

Note: If a CSV disk is put into maintenance mode in a Windows failover cluster manager, it is displayed as dedicated in SnapDrive.

Best Practice

In SnapDrive for Windows, NetApp recommends that you create a CSV from the node that owns the available cluster storage group. Use the `CLUSTER GROUP` command or the `Get-Cluster Group` cmdlet to identify the node that owns the available storage group before you create a CSV disk.

Asymmetric Clustering

Asymmetric clustering is a feature used to create a shared disk or CSV only among a few nodes in a cluster.

Note: SnapDrive 7.x for Windows does not support this feature.

BitLocker Encryption

BitLocker is a data protection feature that has been around since Windows Server 2008 R2. This feature is now available in Windows Server 2012 and Windows Server 2012 R2 with additional functionality. New BitLocker functionality allows you to encrypt cluster-shared SAN volumes. For more information about BitLocker configuration, refer to the Microsoft TechNet article [BitLocker Overview](#).

Note: SnapDrive 7.x for Windows supports BitLocker functionality for CSVs provisioned through SnapDrive.

Virtual Hard Disk Format

Windows Server 2012 and Windows Server 2012 R2 have a new virtual hard disk format (VHDX). Unlike the previous VHD format, this format supports up to 64TB in size. In addition, VHDX has a 4kB logical sector size, which increases the performance of applications that are designed for 4kB sector sizes.

SnapDrive 7.x for Windows supports VHDX. 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: NetApp recommends that you create LUNs smaller than 16TB to host VMs.

Virtual Fibre Channel

Windows Server 2012 and Windows Server 2012 R2 support the provisioning of storage to guest VMs through vFC. Features such as live migration and MPIO are also supported. vFC requires an NPIV-enabled FC host bus adapter (HBA). At most, four FC ports are supported. If the host system is configured with multiple FC ports and presented to the VM, then MPIO must be installed in the VM to enable multipathing. Also, pass-through disks cannot be provisioned to the host if MPIO is being used on the host because pass-through disks do not support MPIO.

Best Practice

If you are using multiple FC initiators on the host, NetApp recommends installing Data ONTAP DSM in the guest VMs to enable multipathing in the guest.

SnapDrive supports provisioning of storage through vFC by using Windows Server 2008 R2 SP1 and Windows Server 2012 and Windows Server 2012 R2 guest VMs. After SnapDrive is installed in the guest VM, SnapDrive recognizes the vFC initiator and enables storage provisioning. Up to 256 LUNs can be provisioned per vFC port presented to the VM. All the SnapDrive features, such as Snapshot and restore, are supported on a vFC LUN.

Note: For information about FC HBAs, refer to the [NetApp Interoperability Matrix](#).

For more information about vFC, refer to the Microsoft TechNet article [Hyper-V Virtual Fibre Channel Overview](#). You can also refer to the Appendix in this technical report for more information about how to use vFC on NetApp storage systems.

Group Managed Service Accounts

Group Managed Service Accounts (gMSA) is a password management solution for service accounts in Windows Server 2012 and Windows Server 2012 R2 environments. It is an improvised version of Managed Service Accounts, which was introduced in Windows Server 2008 R2. Unlike Managed Service Accounts, in gMSA, the password is generated and maintained by the Key Distribution Service on Windows Server 2012 and Windows Server 2012 R2 domain controllers, thereby allowing multiple hosts to use gMSA. Member servers that want to use the gMSA simply query the domain controller for the current password, thereby eliminating the need for an administrator to manually administer passwords for these accounts.

20.4 Windows Server 2012 and Windows Server 2012 R2 Feature Not Supported in SnapDrive

NetApp Data ONTAP, SnapDrive 7.x for Windows, and the NetApp SnapManager suite of products do not support the following feature for Windows Server 2012 and Windows Server 2012 R2:

- Windows Server 2012 and Windows Server 2012 R2 native thin provisioning

20.5 Windows Server 2012 Virtual Machine Support for ESXi Environments

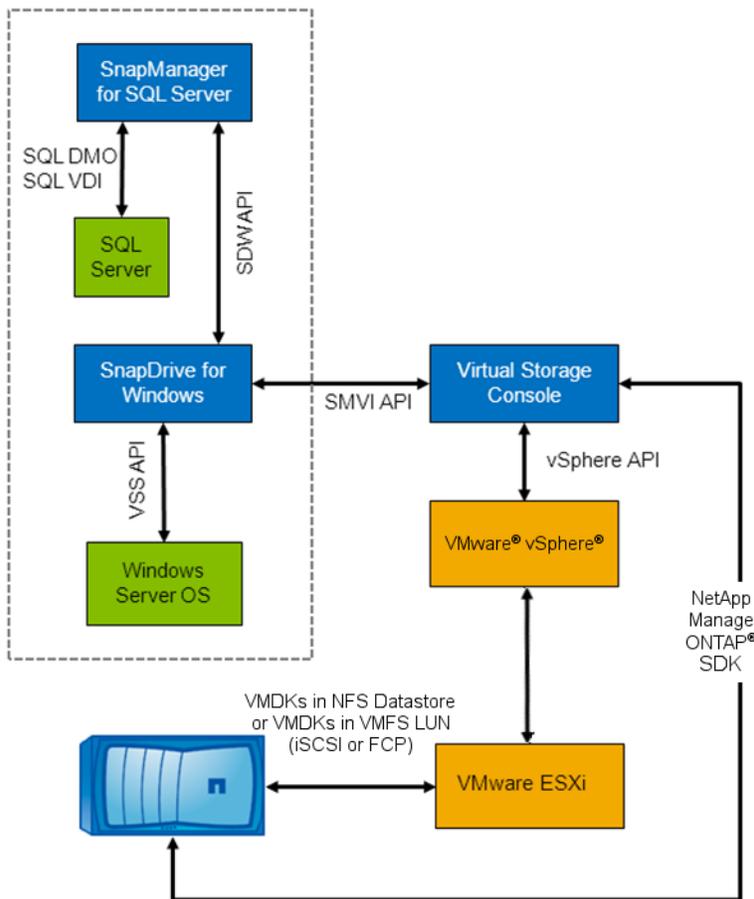
VMware ESXi™ 5.0 U3, ESXi 5.1 U2, and ESXi 5.5 U2 support Windows Server 2012 VMs.

21 VMware Support

SnapDrive for Windows supports provisioning of RDMs on VMFS datastores. Also, VMDKs hosted on VMFS and NFS datastores are supported for both Data ONTAP 7-Mode and clustered Data ONTAP systems.

SnapDrive provides LUN provisioning and file system–consistent backups and recovery by leveraging NetApp storage array Snapshot copies for VMs that are hosted in a VMware environment.

Figure 10) SQL Server database hosted on VMDKs by using SnapDrive and NetApp VSC.



SnapDrive integrates with VSC to identify and list all the VMDKs created in VMware vCenter™. This can be achieved by adding the VMware vCenter or ESX credentials as well as the VSC server details in SnapDrive. After the VMDK is listed in SnapDrive, all operations such as Snapshot creation; Snapshot deletion; mounting, remote mounting, disconnecting, listing, renaming, and restoring of Snapshot copies; space reclamation on NFS datastores; and updating the SnapMirror destination are supported.

When SnapManager products such as SnapManager for SQL Server request a backup operation (for example, SnapManager for SQL Server requests that a Snapshot copy be created), this request is passed to SnapDrive. SnapDrive determines whether the disk for which the Snapshot operation is

requested is a VMDK. If the disk is a VMDK, SnapDrive makes the appropriate vDisk application programming interface (API) call to VSC, and the operation is performed by VSC.

After the backup operations, SnapDrive assumes that all Snapshot backups returned by VSC are application consistent.

Note: Space reclamation is not supported by SnapDrive for VMDKs hosted on VMFS datastores.

Note: Installing VSC 4.1 or later supports clustered Data ONTAP.

Note: IPv6 is not supported with VSC.

Note: VSC does not support VMDKs on LUNs hosted on qtrees.

Best Practices

- SnapDrive must be configured to communicate with the VMware vCenter Server™ to authenticate for web services.
- If vCenter Server is inaccessible, SnapDrive should be reconfigured to communicate with the ESX server by using the SnapDrive Microsoft Management Console or sdcli.
- To enable proper VMware vMotion® functionality, the VMFS, NFS, or RDM datastore hosted on NetApp storage should be shared between the ESX hosts.

22 Performance Tuning

To improve the performance of the SnapDrive system, follow these best practices:

- Add exclusions for SnapDrive service in the antivirus application agent ([NetApp KB ID 2013167](#)) installed on the host server system.
- Configure the Windows firewall to allow SnapDrive services during SnapDrive installation for Windows communications.
- If an SVM has multiple management LIFs, for faster communication, set only one IP as the preferred IP address in the transport protocol settings.
- If the host system has hotfix 2729452 installed, the system might experience a significant delay in system boot time. For information about how to resolve this issue, refer to the [Microsoft Support](#) website.

Note: This applies to Windows Server 2008 R2 and Windows Server 2012 systems only.

Appendix

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

Consider three nodes named node1, node2, and node3, respectively, which are part of a cluster called W2012CLUS. 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.
- The node1 dedicated disk has drive letter E:\.
- The node2 dedicated disk has drive letter F:\.

- A new node, node3, is added to the cluster W2012CLUS through the Add Node wizard. The wizard presents the option to add all eligible storage in the cluster to the user; this option is selected by default.

Behavior

- Node3 is added to the cluster W2012CLUS.
- The node1 dedicated disk (E:\) is converted as follows:
 - The dedicated disk is converted to a clustered disk.
 - A cluster resource is created for the disk that is cluster disk 1.
 - The cluster resource is part of the available storage group.
 - The disk is online and accessible.
- The node2 dedicated disk (F:\) is converted as follows:
 - The dedicated disk is converted to a clustered disk.
 - A cluster resource is created for the disk that is cluster disk 2.
 - The cluster resource is part of the available storage group owned by node1.
 - The disk is offline and not accessible by SnapDrive.
- The dedicated disk of node2, that is, F:\, is not seen by either node1 SnapDrive or node2 SnapDrive:
 - In node2, go to Disk Management and locate the dedicated disk. The disk is offline. A mouse hover displays the tool tip Reserved (the disk is offline because of policy set by an administrator).
 - If you try to bring the disk online from the Disk Management console, an error message is displayed: 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.
 - SnapDrive displays 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 cluster disk 2 online, complete the following steps:

1. Add the iSCSI initiator of node1 to the LUN by using NetApp OnCommand System Manager or Cluster Element Manager.
2. Bring the associated cluster resource cluster disk 2 online in Failover Cluster Manager.
3. Refresh the disks in the SnapDrive MMC console to see cluster disk 2 in SnapDrive.
4. Repeat the same operation for cluster disk 1 to see the disk in SnapDrive.

Version History

Version	Date	Document Version History
Version 1.0	August 2013	Initial release
Version 2.0	December 2014	Minor updates

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

Copyright Information

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

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

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

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

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

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

Trademark Information

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

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. TR-4228-1214