



Technical Report

NetApp Disaster Recovery Solution for Microsoft Hyper-V on Clustered Data ONTAP

Vinith Menon, NetApp

October 2014 | TR-4343

Abstract

This technical report describes disaster recovery options for Microsoft® Hyper-V® 3 on Windows Server® 2012 R2 and provides two scenarios for developing a disaster recovery solution for NetApp® storage. It is intended to be a solution planning guide, providing Hyper-V infrastructure administrators and planners with key architecture components that enable them to successfully design a disaster recovery solution. This technical report discusses the use of NetApp SnapMirror® technology with both clustered and standalone Hyper-V Server configurations.

TABLE OF CONTENTS

1	Introduction	3
1.1	Purpose and Scope	3
1.2	Target Audience	3
2	Disaster Recovery, High Availability, and Business Continuity Overview	3
3	Hyper-V Disaster Recovery Options	4
4	NetApp DR Solution for Microsoft Hyper-V	5
4.1	NetApp Solution Advantages	6
4.2	NetApp Solution Components	7
4.3	NetApp Disaster Recovery Scenarios	8
5	Best Practices	13
6	Summary	14
	Appendix: Hyper-V DR Script Using PowerShell Toolkit	14

LIST OF TABLES

Table 1)	Business problems and solutions in a resilient architecture.	5
Table 2)	Dataset tiers and achievable RPOs and RTOs.	9

LIST OF FIGURES

Figure 1)	SnapMirror setup	8
Figure 2)	Hyper-V hosts with NetApp SnapMirror in a healthy state.	10
Figure 3)	Hyper-V hosts with NetApp SnapMirror after failover	11
Figure 4)	Site A VMs on clustered Hyper-V hosts with failover configured to Site B.	12
Figure 5)	Disaster recovery.	13

1 Introduction

Virtualization has become an integral part of enterprises. System failures can result in operational and financial losses when customers, employees, partners, and other stakeholders are adversely affected by application unavailability. Companies therefore require fast recovery with little or no data loss. When an outage occurs, it can take days to restore service from slower media such as tape, even when all backup media are readily available and error free. Many organizations rely on Microsoft Hyper-V virtualization and failover capabilities, which complement partner data-replication products to build highly effective disaster-recovery (DR) products that protect mission-critical applications.

1.1 Purpose and Scope

This document provides an overview of generic approaches for structuring a disaster recovery model for Microsoft Hyper-V Server using the NetApp comprehensive suite of hardware and software solutions.

The scenarios presented here are designed to achieve multiple levels of recovery point objectives (RPOs) and recovery time objectives (RTOs).

1.2 Target Audience

The target audience for this technical report includes the following roles:

- IT professionals
- Storage professionals
- Hyper-V infrastructure administrators

To understand the methods and procedures covered in this technical report, the reader should have working knowledge of the following concepts:

- Hyper-V component architecture
- Hyper-V infrastructure administration and management
- Service-level expertise regarding Hyper-V recovery options

The reader should also have working knowledge of the following NetApp products and solutions:

- The NetApp Data ONTAP[®] operating system
- SnapDrive[®] for Windows[®]
- SnapManager[®] for Hyper-V (SMHV); backup and restore procedures
- NetApp SnapMirror technology

2 Disaster Recovery, High Availability, and Business Continuity Overview

Any business or organization can be vulnerable to incidents that prevent it from continuing normal operations. Such incidents can be a flood or fire, a serious computer malfunction, or an information security breach.

When a primary production site goes down, disaster recovery (DR) enables the continued operation of virtual machines (VMs) at a secondary, backup site. A DR solution balances three goals: RPO, RTO, and manageability. Use of a solution that meets these goals requires careful design and implementation.

Organizations must consistently find solutions that not only meet application and data requirements for capacity, performance, and availability, but also deliver a positive return on investment (ROI) and cost-reduction capabilities.

Business continuity (BC) refers to the return of business processes to full capability following a disruption of service. Determining which applications are mission critical—essential to an organization’s functioning in case of a disaster—is one of the first steps in high-availability (HA) and BC planning. After crucial applications and services are identified, it is essential to define RPOs and RTOs for them, balancing cost and acceptable risk. An understanding of the following terms is crucial to architecting an optimal DR solution.

- **Availability.** Generally, availability refers to the proportion of time a system, subsystem, service, or piece of equipment is in a fully operational state.
- **Disaster recovery.** DR is the process of making data, hardware, and software available in order to resume critical business operations following a disaster. A robust DR plan consists of the IT, architecture, and processes necessary to recover mission-critical data from a backup location.
- **High availability.** HA is a system-design protocol and implementation that prevents single points of failure and enables the operational continuity of a system, service, or piece of equipment during a given period of measurement.
- **Recovery point objective (RPO).** The RPO is the maximum amount of data loss, in terms of time, that an organization believes it can endure.
- **Recovery time objective (RTO).** The RTO is the maximum amount of downtime allowable before an application is available to users.
- **Manageability.** Manageability is a subjective measure of IT administrative resources that a particular DR solution requires.
- **Service-level agreement (SLA).** A service-level agreement is a formal, negotiated agreement between a service provider and a user (typically a customer) specifying levels of availability, serviceability, performance, and operation of a system, service, or application.
- **Service levels.** To simplify business-continuity and disaster-recovery planning, it’s common to associate a VM or service with a class. These classes then define such things as backup intervals and data retention requirements. Classification of VM workloads (based on backup and availability characteristics) is based on three levels of service:
 - Gold; tier 1 VMs
 - Silver; tier 2 VMs
 - Bronze; tier 3 VMs

Mixing of applications of different service-level classes in a VM is required in order to maximize the efficiency of the underlying hardware. For more details on effectively designing service levels, refer to [TR-4329: NetApp Private Cloud Capacity Planning Guide](#).

3 Hyper-V Disaster Recovery Options

Hyper-V offers many DR, HA, and BC features that vary in their RPO and RTO. These features vary in their relative complexity and resource requirements. The following HA and BC features are available for Hyper-V Server disaster recovery.

- **Hyper-V Replica.** Hyper-V Replica is a new feature of Windows Server 2012. It allows you to asynchronously replicate a VM from a primary to a replica site. It does not require a failover cluster or shared storage. With the release of Windows Server 2012 R2, the replication interval is user-configurable to every 30 seconds, every 5 minutes, or every 15 minutes.

After replication is configured and enabled, an initial copy of data is sent from the primary location to the replica location. Two options are available to perform the initial replication:

- Send the initial copy of the VM over the network, with no additional work.
- If network bandwidth is limited, you can transfer the initial replica using a USB disk and import the initial copy on the replica server. Using a set of PowerShell® cmdlets for Hyper-V Replica, such as `Start-VMInitialReplication`, you can export the initial replica to a location on NetApp

storage. `Start-VMInitialReplication` assumes that the initial location is an external drive that will be shipped to the replica site, but you can also export to NetApp LUNs or SMB shares in order to complete the initial replication. Once this is done, you can transfer the data to the replica server using `SnapMirror` and import the VM replication files from the replica location using `Import-VMInitialReplication`. NetApp `SnapMirror` technology complements this solution by allowing you to replicate the VMs quickly to the secondary site.

The disaster recovery solution design should be based on the prioritization of the VMs. Decide which machines are the highest priority for quick recovery. Use Hyper-V Replica in conjunction with NetApp `SnapMirror` to move the replicated data from the primary site to the replica site. For lower-priority VMs, use `SnapMirror` and standard NetApp recovery processes.

- **Hyper-V live migration.** Hyper-V live migration moves operational VMs from one physical server to another with no impact on the availability of the VMs to users. With Windows Server 2012, Microsoft introduced Microsoft Offloaded Data Transfer (ODX), also known as copy offload, which allows you to offload the job of copying files between two servers to the storage system. Storage live migration and NetApp copy-offload technology complement this solution by providing a mechanism to perform full-file or subfile copies between two directories residing on a single or multiple remote servers.

ODX complements a high-availability scenario, allowing the system administrator to rapidly move data between systems, reducing the exposure time of the application or the VM in flight and putting less strain on the network. It also is useful in a disaster recovery scenario in which a production network is strained as an entire site struggles to come back online. The copy-offload technology is supported for VMs that reside on Cluster Shared Volumes (CSVs) and also SMB shares. It also supports shared-nothing live migration. Shared-nothing live migration allows the migration of a VM—including storage, memory, and device state—between Hyper-V hosts, without any downtime. ODX enables data transfers within or between ODX-enabled storage servers, without transferring the data through the Windows client. This technology helps to effectively reduce the client/server processor/memory utilization and minimizes network I/O bandwidth.

4 NetApp DR Solution for Microsoft Hyper-V

NetApp BC solutions cost effectively protect critical applications and data and enable rapid recovery in the event of a system, site, or regional outage. NetApp DR solutions help maintain availability across a broad spectrum of RPO and RTO requirements during planned and unplanned downtime. With NetApp BC solutions, you can lower the overall cost of high availability and disaster recovery. Ease of administration also improves operational costs. NetApp recovery solutions are easy to implement and operate, placing minimal demands on IT and network administration resources. This solution integrates smoothly into an existing infrastructure and is operational almost immediately, with no need to change the way existing systems and business operations work.

Table 1 provides a description of some common business problems and how to address them at the primary site in order to provide a resilient architecture.

Table 1) Business problems and solutions in a resilient architecture.

Business Problem	Solution	Description
Single point of failure	Hyper-V failover cluster with Cluster Shared Volumes + NetApp storage	Cluster Shared Volumes in a Hyper-V failover cluster address VM resiliency; NetApp storage clusters address resiliency on the storage. Together, these configurations eliminate single points of failure on the VM instance, server hardware, and storage.
Single point of failure	Hyper-V failover cluster + SMB 3.0 shares + NetApp	SMB 3.0 shares in a Hyper-V failover cluster address VM resiliency; NetApp storage

Business Problem	Solution	Description
	storage	clusters address resiliency on the storage. Combined, they eliminate single points of failure on the VM instance, server hardware, and storage.
Fast backup and recovery	SnapManager for Hyper-V (SMHV)	SMHV automates a complex, manual backup process by creating fast, space-efficient NetApp Snapshot™ copies and providing fast recovery by using the SMHV GUI and PowerShell cmdlets.
Disaster recovery	SnapManager for Hyper-V + SnapDrive for Windows (SDW) + SnapMirror	SnapMirror replicates VM components and SMHV/SDW components to provide faster backups and rapid restores.

4.1 NetApp Solution Advantages

The following key advantages can be leveraged using this solution to provide cross-site resiliency and high availability.

- **Ease of configuration.** The most advantageous aspect of this disaster recovery solution is the ease with which it can be deployed. Using the SMHV GUI and a few PowerShell cmdlets, a robust DR solution can be deployed easily and effectively. This reduces administrative overhead in any environment.
- **Speed and performance.** SMHV backups are based on NetApp Snapshot technology, which minimizes how long the backed-up VM remains in a saved state. SMHV software also supports the use of PowerShell cmdlets that simplify, automate, and speed up the DR process. This easy-to-use management tool lets you leverage NetApp Snapshot copy and replication technology for near-instant backups. Additionally, asynchronous SnapMirror updates are remarkably fast and allow healthy RPOs and RTOs to be achieved.
- **VM backup options.** SMHV provides two flavors of VM backup:
 - Application-consistent backups are generally made of VMs with applications installed on the local operating system and not backed up elsewhere. A typical example of this is an application-consistent backup of a domain controller. The Active Directory (AD) servers must not be recovered from replicas created by unsupported processes because this can create an update sequence number rollback scenario in which an AD server that contains an older version of the AD database is recovered from an unsupported backup method. This can cause the AD processes on the recovered server to be unable to process authentication requests or other functions. When an application-consistent backup is invoked, the Hyper-V writer calls the application Volume Shadow Copy Service (VSS) writer, quiesces the applications and file system inside the VM, and commits any writes before the volume-level Snapshot copy is made.
 - Crash-consistent backup jobs are quick Snapshot copies of VMs in a dataset. The resulting backup copies are similar to the data captures of VMs that crash or are abruptly powered off. Crash-consistent backups make use of only Snapshot copies, and do not provide VSS integration. Crash-consistent backups can be made of applications that are generally backed up by another backup utility. They can effectively back up the VM state for VMs that have application data residing on NetApp LUNs rather than virtual hard disks (for example, databases for various Microsoft applications such as Exchange, SQL Server®, and SharePoint®). Once the VM-level backup is performed by invoking a crash-consistent backup, the SnapManager suite can be used to perform a backup of application data for these applications.
- **Lower costs.** Key NetApp technologies help reduce infrastructure costs. NetApp SnapMirror technology can be used to move data between sites. This is very useful in DR plans because

volumes can be incrementally replicated to an off-site location. NetApp SnapManager for Hyper-V allows you to avoid bottlenecks by offloading backup jobs from the host server to the storage system, thereby speeding up backup, restore, and disaster-recovery tasks in a virtualized Windows environment.

- SnapMirror technology simplifies DR operations and allows customers to lower their RTO while achieving RPO, based on their own SLAs. By replicating only the changed blocks of data after the initial baseline transfer, SnapMirror decreases the replication time across sites, reducing network usage and storage costs. SnapMirror can be deployed with no additional IT resources and supports frequent testing of your disaster recovery plan. SnapMirror can be used to extend protection by replicating backups to a different storage system—either within the data center, to another local data center, or to a remote DR site.
- SnapManager for Hyper-V allows administrators to make backups of online VMs that have Hyper-V integration services installed. Saved-state VM backups are made of VMs that don't have the latest version of integration services. The application-consistent and crash-consistent modes of backup provide varying degrees of recovery capability in between full backup jobs in the event of sudden loss of the primary site. Crash-consistent and application-consistent backup types can be used together in a SMHV dataset, thus making it easy to group a complex IT infrastructure into a manageable backup and disaster-recovery plan.
- **Simplified Hyper-V systemwide DR.** By using the simplified PowerShell scripting and command-line features integrated in SMHV, an administrator can deploy DR plans per VM instance or on multiple VMs. No special or extra steps need be taken to impart DR capability on the VMs.

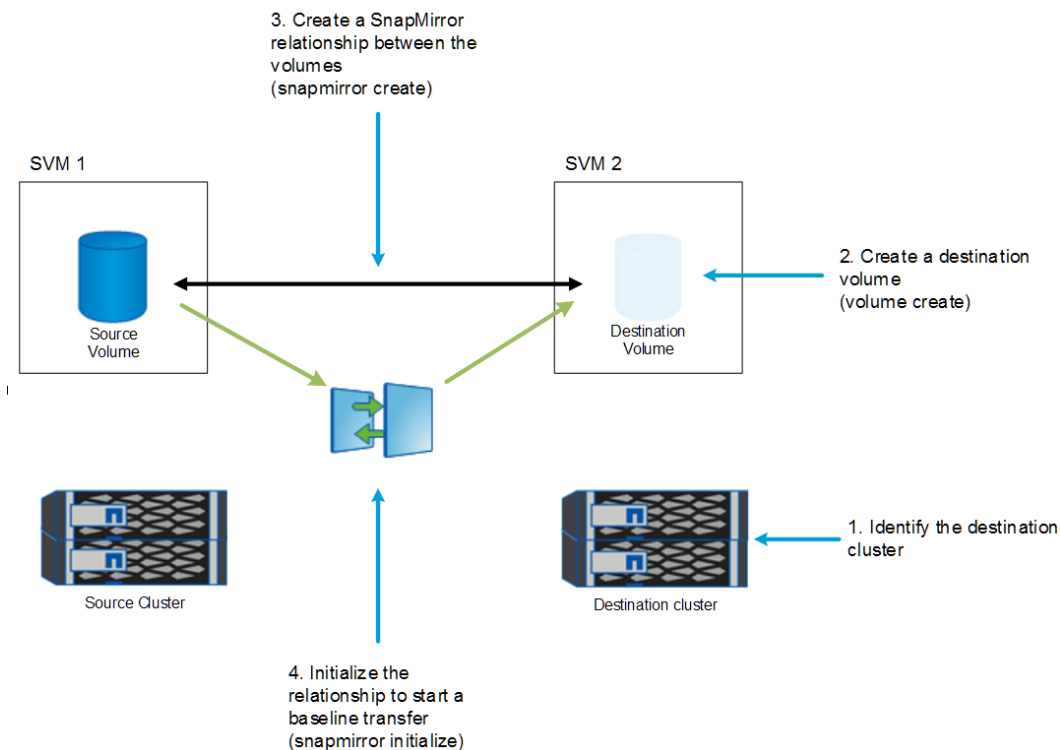
4.2 NetApp Solution Components

This solution integrates the following NetApp software components.

- **SnapDrive for Windows.** NetApp SnapDrive for Windows is an enterprise-class, storage- and data-management solution for Windows Server environments. SnapDrive enables storage and system administrators to quickly and easily manage, map, and migrate data.
- **SnapManager for Hyper-V.** NetApp SnapManager for Hyper-V software automates and simplifies backup and restore operations for Microsoft Windows Server 2012 R2 with Hyper-V environments hosted on NetApp storage. This easy-to-use management tool lets you leverage NetApp Snapshot copies and replication technology for instant, application-consistent backups and cost-effective DR for your entire Microsoft virtual environment. NetApp SnapManager for Hyper-V offloads backups from the host server to the storage system, leading to high operational efficiency.
- **SnapMirror.** NetApp SnapMirror enables efficient site-to-site storage replication, making disaster recovery rapid, reliable, and manageable to suit today's global enterprises. Replicating data at high speeds over LANs and WANs, SnapMirror provides high data availability and fast recovery for mission-critical applications, as well as outstanding storage deduplication and network compression capabilities.

With NetApp SnapMirror technology, disaster recovery can protect the entire data center. Volumes can back up to an off-site location incrementally. SnapMirror for Hyper-V (SMHV) integrates natively with SnapMirror technology, enabling the administrator to easily ship backups to the remote DR site immediately following initiation on the primary site. SnapMirror performs incremental, block-based replication as frequently as the required RPO. The block-level updates reduce bandwidth and time requirements, and data consistency is maintained at the DR site. SMHV enables restores of the VM to a point in time before data corruption occurred. Figure 1 illustrates the procedure for initializing a SnapMirror relationship.

Figure 1) SnapMirror setup.



The first and most important step is to create a one-time baseline transfer of the entire dataset. This is required before incremental updates can be performed. This operation includes the creation of a Snapshot copy at the source and the transfer of all of the data blocks referenced by it to the destination file system. After the initialization is complete, scheduled or manually triggered updates can occur. Each update transfers only the new and changed blocks from the source to the destination file system. This operation includes creating a Snapshot copy at the source volume, comparing it with the baseline copy, and transferring only the changed blocks to the destination volume. The new copy becomes the baseline copy for the next update.

Because the replication is periodic, SnapMirror is able to consolidate the changed blocks and conserve network bandwidth. The impact on write throughput and write latency is minimal.

4.3 NetApp Disaster Recovery Scenarios

The primary objective of this DR solution is to achieve the highest degree of operational continuity at the primary site while eliminating any single point of failure.

Two scenarios for achieving optimal levels of RPOs and RTOs are described in this section. When designing a disaster recovery solution for Microsoft Hyper-V, it is important to review your current SLAs to derive appropriate RPOs and RTOs.

Scenario 1) Disaster Recovery Using SnapManager for Hyper-V

The disaster recovery solution described in Scenario 1 uses SMHV. In NetApp SMHV terminology, VMs can be grouped into datasets comprised of tier 1, tier 2, and tier 3 VMs. Tiering VMs is a technique used to define backup and availability characteristics. The duration of backup and restore operations is affected by the number of VMs present. The level of service determines how many VMs you can place in a Cluster Shared Volume or on an SMB 3.0 share. Smaller Cluster Shared Volumes or SMB shares restore more quickly than larger ones. Similarly, large numbers of Cluster Shared Volumes on a node increase the

RTO. In most cases, a combination of VM tiers is used in order to maximize the efficiency of your hardware.

Table 2 shows the RPO and RTO achieved by grouping the VMs into different dataset tiers, based on priority.

Table 2) Dataset tiers and achievable RPOs and RTOs.

Dataset Tier	Virtual Machines per Dataset	RPO (Minutes)	RTO (Minutes)
Tier 1	50	5	15
Tier 2	100	10	20
Tier 3	200	20	30

Note: The figures in Table 2 are derived from NetApp testing and will vary depending on deployment variables.

This architecture provides high availability and disaster recovery for VMs running on Microsoft Hyper-V in a primary site for production and a secondary, DR site for recovery of VMs. Figure 2 shows the architecture prior to disaster recovery.

Figure 2) Hyper-V hosts with NetApp SnapMirror in a healthy state.

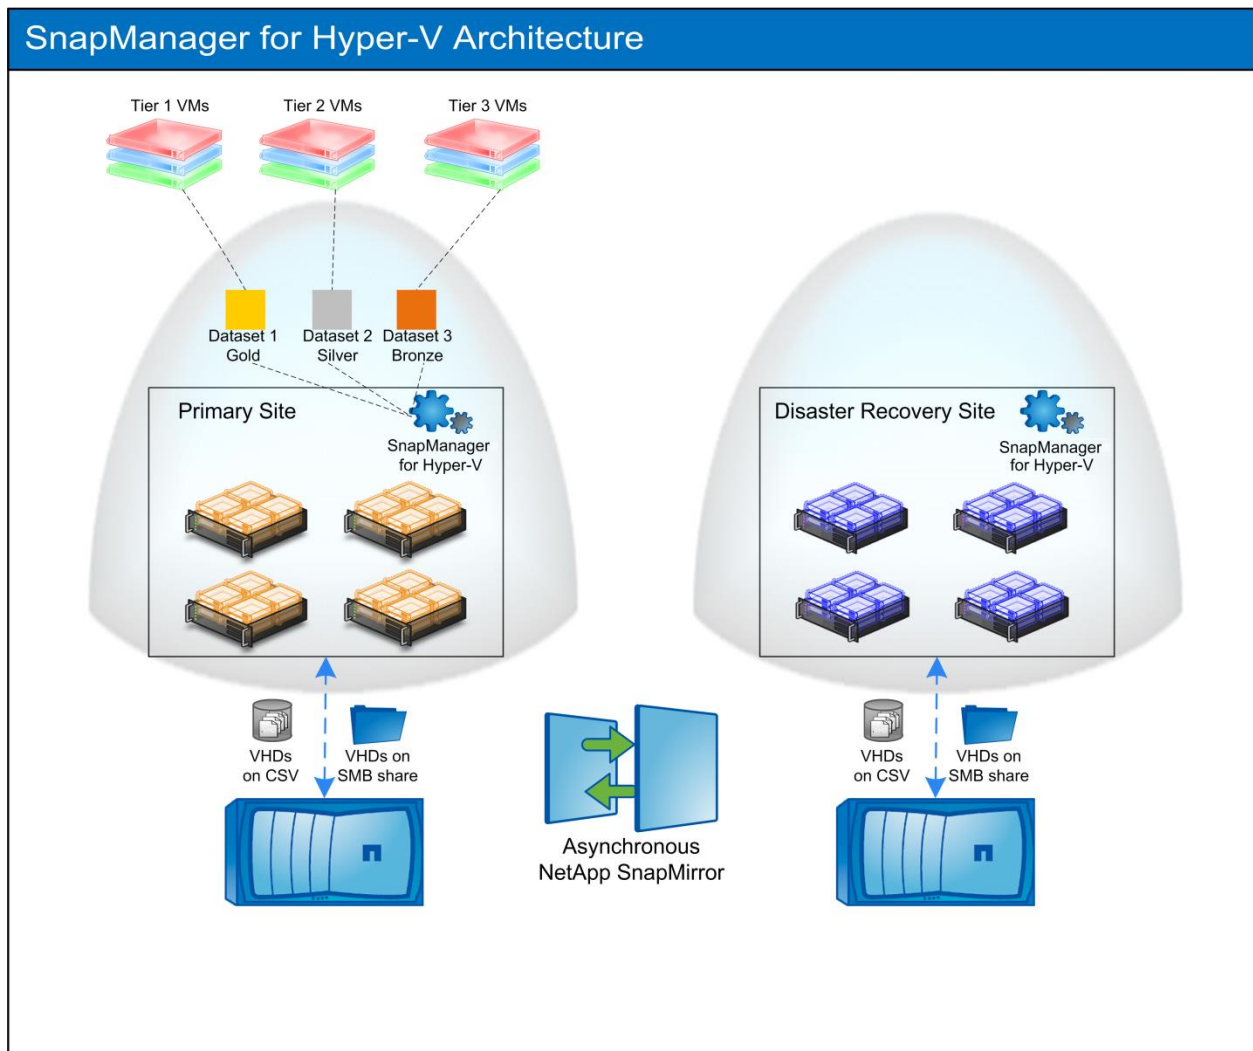



Figure 3 illustrates the architecture after failover.

Figure 3) Hyper-V hosts with NetApp SnapMirror after failover.



Disaster Recovery Configuration Procedure

1. Install NetApp host utilities for Windows, SnapDrive for Windows, and SnapManager for Hyper-V on the Hyper-V cluster on the disaster recovery site.
2. Determine that the backup server is connected properly using iSCSI or Fibre Channel (FC) to the SnapMirror destination storage.
3. Use SnapDrive for Windows to connect to the LUNs at the SnapMirror destination. Use the same drive letters as the original server or drive letters for the secondary server configuration. SnapDrive for Windows will automatically break the SnapMirror relationship.
4. Once the LUNs are mapped to the Hyper-V host, the VMs can be recovered from the most recent backups using the SMHV Windows PowerShell cmdlets, which will provide the required RPO.

Scenario 2) Disaster Recovery Using Windows PowerShell

The solution architecture for this scenario is the same as that of Scenario 1, above. Developed by NetApp, the Data ONTAP PowerShell Toolkit (PSTK) is a set of cmdlets for Windows PowerShell that includes many of the APIs available in the NetApp Manage ONTAP® Software Developer Kit (SDK). The cmdlets enable you to tap into the power of NetApp APIs to easily perform a variety of NetApp specific tasks. Over 1,400 cmdlets are available, covering many aspects of NetApp storage system operations.

You can automate and create a complete Hyper-V disaster recovery solution using PSTK cmdlets and Windows Server 2012 Hyper-V 3.

The scripting workflow enables the disaster recovery of VMs residing on Cluster Shared Volumes (CSVs).

Scenario 2 features a set of highly available VMs at Site A running on clustered Hyper-V hosts serving mission-critical applications and that have Cluster Shared Volumes hosted on a storage virtual machine, SVM 1.

When the VMs experience complete site failure (both host and storage), the storage administrator can instantaneously fail over all resources to SVM 2, hosted at Site B.

Using PowerShell cmdlets, the administrator initiates Hyper-V disaster recovery. Consistent VMs can then serve mission-critical applications from a replicated or mirrored Site B nearly instantaneously and with minimal downtime.

SnapMirror relationships are established between volumes in the SVMs hosted at Site A and Site B.

Figure 4) Site A VMs on clustered Hyper-V hosts with failover configured to Site B.




Figure 4 illustrates the disaster recovery architecture. Site A consists of a Hyper-V host that's cluster connected to NetApp storage volume 1 (Vol1) on SVM 1. The failover site, Site B, consists of a Hyper-V host cluster connected to storage volume 2 (Vol2) on SVM 2. A SnapMirror relationship is established between Vol1 and Vol2 in SVM 1 and SVM 2.

The VM configuration files for the VMs hosted on Hyper-V host cluster A are exported to a folder named `Exported VMs` within the CSV of Hyper-V cluster A, which is created as part of the scripting process.

After the VM configuration files are exported, perform a SnapMirror update from Vol1 on SVM 1 to the destination volume, Vol2 on SVM 2. This results in the replication of the latest set of VM configuration files to Vol2 on Site B.

Disaster Recovery Procedure

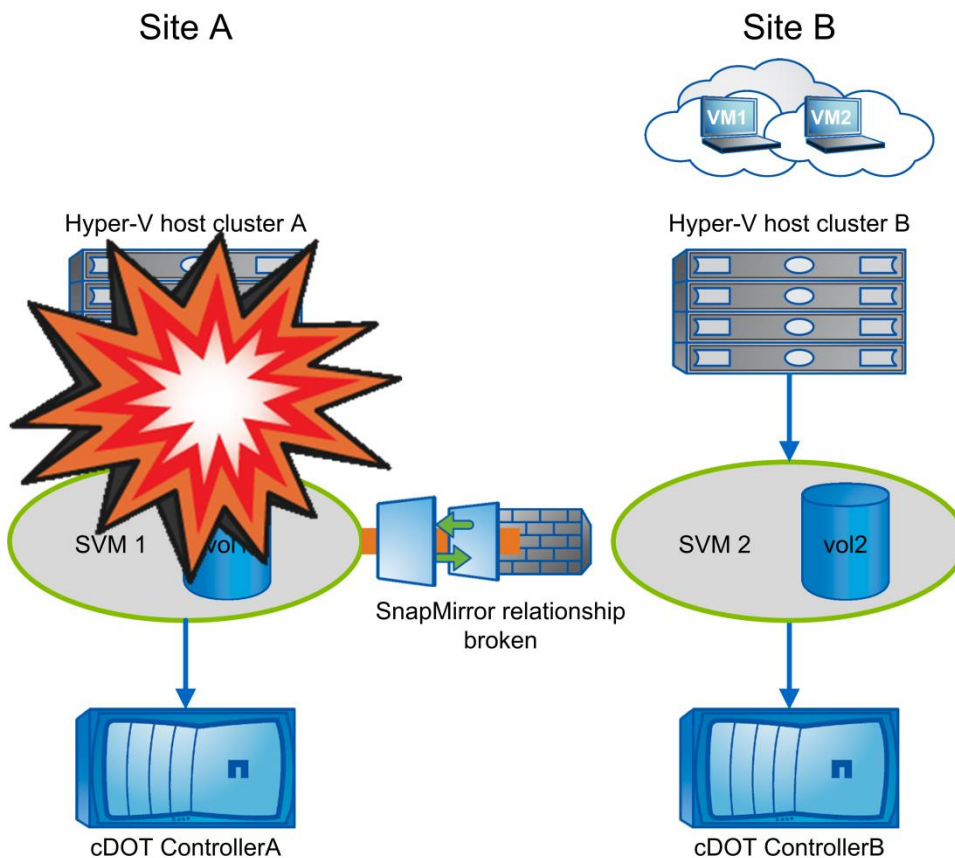
After all of the VM configuration files have been exported successfully, when a disaster scenario occurs and Site A goes down, recovery is performed by completing the following steps:

1. Connect to the controller on the secondary site.
2. Break the SnapMirror relationship.
3. Map the LUNs in the SnapMirror volume to the initiator group (igroup) for the Hyper-V servers on the secondary site.
4. Once the LUNs are mapped to the Hyper-V cluster, make these disks online.
5. Using the failover-cluster PowerShell cmdlets, add the disks to available storage and convert them to CSVs.
6. Import the virtual machines in the CSV to the Hyper-V manager, make them highly available, and then add them to the cluster.
7. Turn on the VMs.

Note: PowerShell scripts are provided in the Appendix.

Figure 5 illustrates disaster recovery.

Figure 5) Disaster recovery.



5 Best Practices

For best practices regarding NetApp technologies, such as SnapManager for Hyper-V, SnapDrive for Windows, and SnapMirror technology, refer to the following documents:

- [TR-4226: NetApp SnapManager 2.0 for Hyper-V on Clustered Data ONTAP 8.2 Best Practices Guide](#)
- [TR-4228: SnapDrive 7.0 for Windows for Clustered Data ONTAP 8.2 Best Practices Guide for SAN Environment](#)
- [TR-4218: SnapDrive 7.0 for Windows SMB 3.0 Best Practices and Deployment Guide](#)
- [TR-4015: SnapMirror Configuration and Best Practices Guide for Clustered Data ONTAP](#)

6 Summary

As a building block of a virtualized infrastructure, Microsoft Hyper-V supports the high availability of mission-critical applications. If these hypervisor-hosted apps become unavailable, an organization's operational productivity can be crippled. NetApp offers proven tools for data protection and disaster recovery on Microsoft Hyper-V servers. SnapManager for Hyper-V backup and restore capabilities provide, along with SnapDrive SnapMirror technologies, a robust solution for protecting and recovering your VMs while meeting stringent RPOs and RTOs that are based on your unique business requirements.

Appendix: Hyper-V DR Script Using PowerShell Toolkit

The following PowerShell script can be tuned to a customer's specific needs when implementing an automated DR solution.

Start by exporting the VM configuration files for the VMs hosted on the Hyper-V cluster A. For the current scenario, export the VMs to a new folder named `Exported VMs` within the CSV of Hyper-V cluster A (which also gets created as part of the scripting process).

Run the following script on any node on the primary Hyper-V cluster on Site A to complete the export process.

```
#### Extracting the Highly Available VM's present in Hyper-V Cluster and storing their
configuration files in CSV ####

# Import the Failover Cluster module as it contains the relevant cmdlets which are needed to
export the VM configuration files.
Import-Module failoverclusters

# Get a list of all cluster shared volumes and extract their respective friendly volume names.
$frndvol = Get-ClusterSharedVolume | select -ExpandProperty sharedvolumeinfo | select -
ExpandProperty friendlyvolumename

# Initialize an array which would be used later to hold the configuration data.
$a = @()

# Extract list of all Highly Available VM's and copy and paste their vm configuration files into
"Exported VM's" folder.
foreach ($frnd in $frndvol)

{

# Extract the VMname, VMid and VM configuration location from all Hyper-V cluster nodes.
$vmdetails = Get-ClusterNode | % {Get-VM -ComputerName $_.name} | select
vmname, vmid, ConfigurationLocation

# Force create an ExportedVM's directory to store the VM configuration files
New-Item -Path "$frnd\ExportedVMs" -ItemType directory -Force

# For each VM copy the VM configuration files to the ExportedVM's directory
foreach ($vm in $vmdetails)

{

$location = $vm.ConfigurationLocation
```

```

$vmname = $vm.vmname
$vmid = $vm.vmid
Copy-Item -Path "$location\Virtual Machines\$vmid.xml" -Destination "$frnd\ExportedVMs"
$a += Get-ChildItem "$frnd\ExportedVMs" | select name
}
}

# Extracting exact unique location of vmconfigfiles
$vmconfigfiles = $a | select * -Unique

####
####

```

Once the VM configuration files are copied, invoke a SnapMirror update from vol1 on SVM 1 to the destination volume vol2 on SVM 2. This replicates the latest set of VM configuration files to vol2 at Site B.

```

# We would start by connecting to the c-mode controller on the secondary site
Connect-NcController $cmodesecsecondariesite -Credential $cred2

# Next we would invoke a snapmirror update from the source vserver svm1 volume vol1 to the
destination vserver svm2 vol2 so that we get the latest data
Invoke-NcSnapmirrorUpdate -DestinationVserver $cmodesecsecondaryvserver -DestinationVolume
$destinationvolume -SourceVserver $cmodeprimaryvserver -SourceVolume $sourcevolume

```

All VM configuration files should by now have been exported. Envision a scenario in which some form of disaster strikes and Site A goes down. The following PowerShell-based steps are run to recover VMs on Site B.

Note: These steps can be run from any node on Hyper-V cluster B.

```

#####
#####
#####

# We would start by connecting to the c-mode controller on the secondary site
Connect-NcController $cmodesecsecondariesite -Credential $cred2

# As we have the latest replicated information from our earlier snapmirror update cmdlet,we would
invoke a snapmirror break so that we can start using vole for our SiteB
Invoke-NcSnapmirrorBreak -DestinationVserver $cmodesecsecondaryvserver -DestinationVolume
$destinationvolume -SourceVserver $cmodeprimaryvserver -SourceVolume $sourcevolume

# Next we will get the igroup information for the hyper-v servers in secondary site
Get-NcIgroup

# Next we will get the luns which we would map and attach to the Hyper-V clusters which would be
eventually converted to a cluster shared volume.
$lunpath = Get-NcLun | ?{$_path -match $destinationvolume} | select -ExpandProperty path

# Next we will map this lun to the secondary site igroup.
Add-NcLunMap -Path $lunpath -InitiatorGroup (Get-NcIgroup | select -ExpandProperty name |
?{$_name -match $igroup}) -VserverContext $cmodesecsecondaryvserver

# Next we need to run a host disk rescan on the host side and discover the disks

Start-NcHostDiskRescan; Wait-NcHostDisk -SettlingTime 5000

```

```

# As the disks are in offline state, lets gets all the disks which are offline and make them
online
Get-Disk | ? IsOffline | Set-Disk -IsOffline:$false

# Next we would add these disks to available disks and convert them to cluster shared volume
$clusterdiskname = Get-ClusterAvailableDisk | select -ExpandProperty name
Get-ClusterAvailableDisk | Add-ClusterDisk
Get-ClusterResource | ?{$_name -match $clusterdiskname} | Add-ClusterSharedVolume

# Once we have the disk's added as csv's we will import the VM's into Hyper-V
$vmconfigfiles = Get-ChildItem "$frndvol\ExportedVMs" | select -ExpandProperty name

foreach ($vm in $vmconfigfiles)

{

Import-VM -Path "$frndvol\ExportedVMs\$vm"

}

# Next we will make all the Virtual Machines highly available
Get-VM | Add-ClusterVirtualMachineRole

# Our final step will be to turn on all the highly available VM's
Get-ClusterNode | % {Get-VM -ComputerName $_.name} | % { Start-VM -Name $_.name -ComputerName
$_.ComputerName}

####
####
####
####

```

Acknowledgement

Special thanks to the following persons for their contributions to this document:

- Niyaz Mohamed
- Glenn Sizemore

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

Copyright Information

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

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

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

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

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

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

Trademark Information

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

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