Technical Report

NetApp Storage Best Practices for Microsoft Virtualization and NetApp SnapManager for Hyper-V

Santhosh Harihara Rao, NetApp
February 2013 | TR-3702

Abstract

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

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Executive Summary

Server virtualization is a major component of data center virtualization and plays a key role in the virtualization initiative. Microsoft is a lead player in this initiative with its server virtualization solutions. This technical report provides detailed guidance on how to architect and implement Microsoft Server virtualization solutions on NetApp storage. It consists of two parts: Chapters 1 through 11 provide details on the best integration points for each of the key enabling NetApp technologies and explain how each technology concept plays a critical role and complements the others to work together as an integrated NetApp solution for Microsoft Server virtualization. Chapter 12 and following sections describe the use of and best practices for SnapManager® for Hyper-V, a NetApp tool that uses the NetApp Snapshot™ technology for backup and recovery of virtual machines in a Hyper-V environment.

NetApp has been on the forefront of solving complex business problems with its innovative technology breakthroughs and end-to-end solutions approach. This technical report is not intended to be a definitive implementation or solutions guide. Expertise might be required to solve specific deployments. Contact your local NetApp sales representative to speak with one of our Microsoft Hyper-V solutions experts. We are dedicated to helping you transform your data center to help your business go further, faster.

2 Server Configuration

2.1 Microsoft Hyper-V R2

Microsoft Windows Server® 2008 R2 allows you to optimize your server hardware investments by providing a scalable, reliable, and secure virtualization platform to consolidate multiple system roles as separate virtual machines (VMs) running on a single physical machine. Typical implementation scenarios for Microsoft Hyper-V virtualization include private and public clouds, production server consolidation, test and development, and business continuity management. For a full list of supported guest OSs on Hyper-V, see Virtualization with Hyper-V: Supported Guest Operating Systems.

To install Hyper-V, NetApp recommends that you follow Microsoft's recommendations.

To enable Hyper-V within a full installation of Windows Server 2008 R2, refer to these resources:

- Install the Hyper-V Role on a Full Installation of Windows Server 2008 on Microsoft TechNet
- TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide

To enable Hyper-V within a server core installation of Windows Server 2008 R2, refer to these resources:

- Install the Hyper-V Role on a Server Core Installation of Windows Server 2008 on Microsoft TechNet
- TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide

To configure Hyper-V Server 2008 R2, refer to the Microsoft Hyper-V Server 2008 Configuration Guide.

After Hyper-V is installed, there are many additional considerations to be made, from configuring the virtual networks to understanding the options for further configuration of Hyper-V. NetApp recommends that you follow Microsoft's recommendations whenever possible; abundant information regarding Hyper-V can be found by searching Microsoft's Web sites.

2.2 Microsoft System Center Virtual Machine Manager Configuration

Microsoft System Center Virtual Machine Manager (SCVMM) is a component of the Microsoft System Center suite of products. SCVMM enables management of heterogeneous environments, both physical and virtual, through a single interface. SCVMM 2008 R2 supports the management of Hyper-V hosts and VMs, as well as VMware® hosts and VMs, and offers a number of other important virtualization tools as well. In addition, SCVMM can be configured to integrate with Systems Center Operations Manager 2007 to provide monitoring information on the servers and VMs that it manages.

Consult the Infrastructure Planning and Design Guide for System Center Virtual Machine Manager 2008 R2 (IPD) before deploying SCVMM 2008 R2. It is important to review the IPD as thoroughly as possible because there are many key decisions required before actually deploying SCVMM 2008 R2, including choosing whether to use a storage area network (SAN) with SCVMM 2008 R2.

Figure 1 shows the decisions and tasks performed before deploying.

Figure 1) SCVMM deployment decision chart.

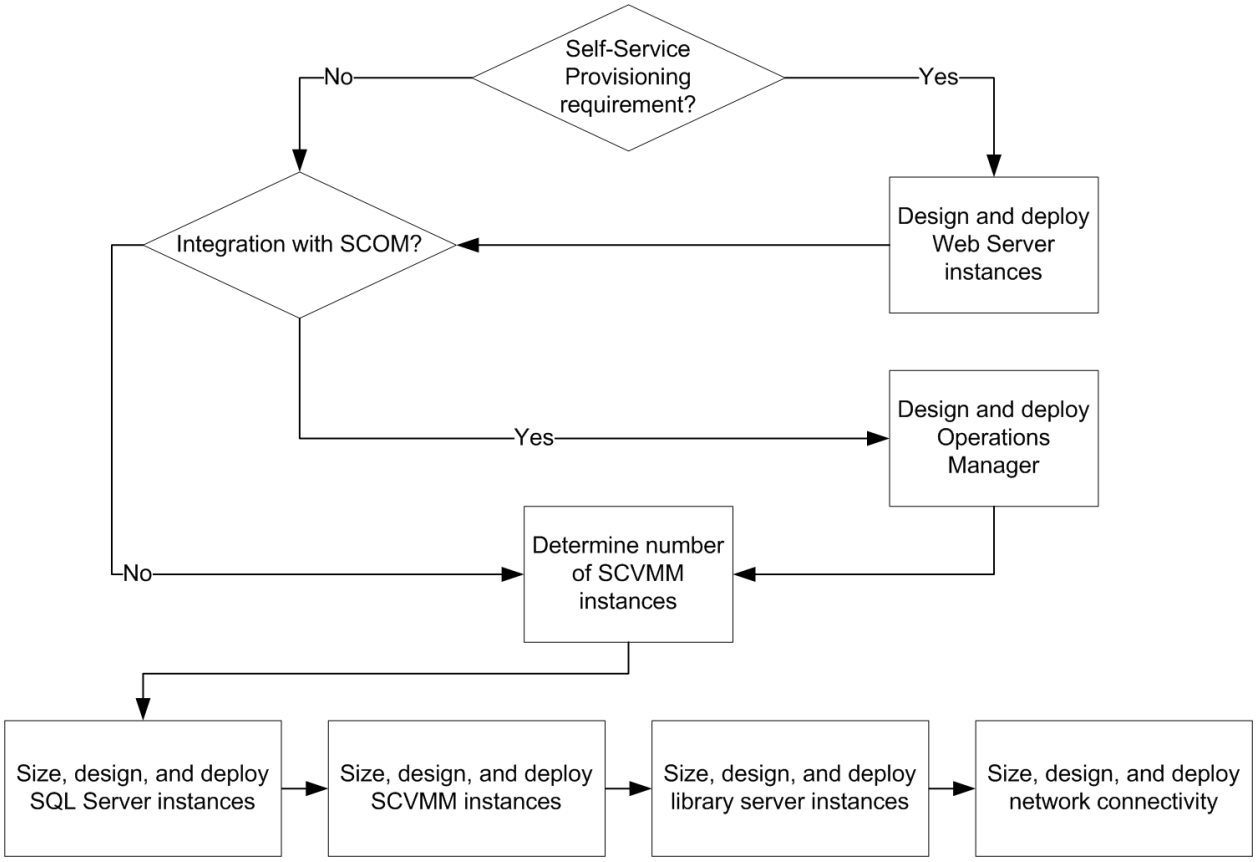

In addition, to support specific features of SCVMM 2008 R2 involving a SAN configuration, there are specific configuration steps that must be followed. For more information, see Configuring a SAN Environment for VMM on Microsoft TechNet.

Best Practice

NetApp recommends that you configure a single FlexVol® volume and logical unit numbers (LUNs) as described later in this technical report. A key point is to configure the appropriate amount of storage for the library components to use as the storage location for the SCVMM library. This minimizes the resources required from the local server to support the sometimes disk-intensive operations associated with the SCVMM library and takes advantage of multiple NetApp storage features to more efficiently manage the data associated with the SCVMM library.

For more information on the installation of SCVMM 2008 R2, refer to the “SCVMM 2008 R2 Installation” section in TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide and follow the instructions to install the SCVMM server components and the SCVMM administrator console. In addition, see New Installation of VMM on Microsoft TechNet.

3 Network Configuration

In an enterprise production environment, availability of applications and data is critical. This is why it is important to plan not just the initial deployment of a new network infrastructure or an upgrade to an existing one, but also the times when you make significant additions to your network infrastructure. Implementing a server virtualization environment adds a significant number of new ports to the network, as the Hyper-V servers often have four or more physical networks adapters installed, and the virtual machines also add additional ports to the network. Although the VM network adapters are virtualized, as well as the virtual switches inside the hypervisor to which they connect, they also must be managed. All the new ports and new virtual networking options often account for a large addition to an existing network and should be planned accordingly.

3.1 Hyper-V Server Networking Considerations

This section discusses the following aspects of Hyper-V server networking:

- Physical network adapters
- Virtual networks
- Network feature support
- Network naming standard
- Network adapter binding order and metric values

Physical Network Adapters

Most Hyper-V servers have four or more physical network adapters installed to handle Hyper-V management, virtual machine connectivity, IP storage connectivity, Windows® failover cluster (WFC) or WFC heartbeat communication, live migration communication, and cluster shared volume (CSV) communication. Smaller environments require a minimum of 2 to 3 network adapters, whereas larger environments require at least 4 to 5 network adapters. Why do we need multiple physical network adapters?

- **Hyper-V management.** Microsoft has constantly recommended that the Hyper-V parent partition, also known as management operating system (MOS), have a dedicated physical network adapter for management of the Hyper-V server as a best practice. Communication is necessary to manage the Hyper-V server and any VMs it hosts remotely, from another system or from SCVMM; therefore, you should consider using network interface card (NIC) teaming to provide redundancy. For more information, see Configuring Virtual Networks on Microsoft TechNet.
- **Virtual machines.** VMs can communicate over external, internal, and private virtual networks that are implemented through the Hyper-V parent partition. Each external virtual switch must map to an individual physical network adapter or logical network adapter from the result of NIC teaming, which is discussed later in this document. To provide redundancy in a production environment, you can assign an external virtual switch to a network team or use multiple external virtual switches; for both configurations, a minimum of two physical network adapters would be needed to provide redundancy. For more information, see Configuring Virtual Networks on Microsoft TechNet.
- **IP storage.** Microsoft recommends that IP storage communication be separate from virtual machine and cluster communications as a best practice, which NetApp supports. Therefore, a minimum of one physical network adapter is required to support iSCSI communication from the Hyper-V parent partition. If you want to use multipathing or multipath input/output (MPIO) from the Hyper-V parent partition, then a minimum of two physical network adapters is required. If you are enabling Windows failover clustering for Hyper-V, maintaining separation of IP storage traffic becomes a requirement for configuration before validating a failover cluster. For more information, see Hyper-V: Using Hyper-V and Failover Clustering and Hyper-V: Using Live Migration with Cluster Shared Volumes in Windows Server 2008 R2 on Microsoft TechNet.

- **Windows failover cluster private.** If you create a Windows failover cluster for Hyper-V, it requires a cluster private network and therefore might require a dedicated physical network adapter. In previous versions of Windows Server, this was used primarily for the cluster heartbeat communications, but with R2 it is also used for cluster shared volumes or live migration communications (see "Live Migration" and "Cluster shared volumes" in this section). For more information, see Hyper-V: Using Hyper-V and Failover Clustering and Hyper-V: Using Live Migration with Cluster Shared Volumes in Windows Server 2008 R2 on Microsoft TechNet.
- **Live migration.** This is a new feature for Windows Server 2008 R2 and does not apply to previous versions of Hyper-V before R2. When live migration of virtual machines is taking place, the communications for facilitating this traverse the network. Microsoft recommends configuring a dedicated physical network adapter for only live migration traffic within the Failover Cluster Manager MMC or using Windows PowerShell™. For more information, see Hyper-V Live Migration Overview and Architecture in the Microsoft Download Center and Hyper-V: Using Live Migration with Cluster Shared Volumes in Windows Server 2008 R2 on Microsoft TechNet.
- **Cluster shared volumes.** CSVs are also a new feature for R2; therefore this section does not apply to previous versions of Hyper-V before R2. When CSVs are enabled within a Windows failover cluster for Hyper-V, there is communication between Hyper-V cluster nodes that are owners and nonowners of a particular CSV, which includes health checks and dynamic I/O redirection. Microsoft recommends a dedicated physical network adapter to make sure there is the necessary bandwidth to support these operations and minimize the event of a failover caused by the inability to support CSV communication between nodes. For more information, see Hyper-V: Using Live Migration with Cluster Shared Volumes in Windows Server 2008 R2 on Microsoft TechNet.

As you can see, the number of physical network adapters recommended per Hyper-V server adds up quickly, especially if you are configuring Windows failover cluster for high availability or using live migration and CSV.

To summarize the recommended number of physical network adapters for specific configurations, see Table 1 through Table 4.

Table 1) Standalone Hyper-V server configuration.

Environment	Protocol	Mgmt	VMs	iSCSI	Cluster	Migration	CSV	Total
Nonproduction	DAS	1	1	N/A	N/A	N/A	N/A	2
	FC	1	1	N/A	N/A	N/A	N/A	2
	iSCSI	1	1	N/A	N/A	N/A	N/A	3
Production	DAS	1	1 or 2	N/A	N/A	N/A	N/A	2 or 3
	FC	1	1 or 2	N/A	N/A	N/A	N/A	2 or 3
	iSCSI	1	1 or 2	2	N/A	N/A	N/A	4 or 5

Table 2) Clustered Hyper-V server configuration.

Environment	Protocol	Mgmt	VMs	iSCSI	Cluster	Migration	CSV	Total
Nonproduction	FC	1	1	N/A	1	N/A	N/A	3
	iSCSI	1	1	1	1	N/A	N/A	4
	FC	1	1 or 2	N/A	1	N/A	N/A	3 or 4
Production	iSCSI	1	1 or 2	2	1	N/A	N/A	5 or 6

Table 3) Clustered Hyper-V servers using live migration configuration.

Environment	Protocol	Mgmt	VMs	iSCSI	Cluster	Migration	CSV	Total
Nonproduction	FC	1	1	N/A	1	1	N/A	4
	iSCSI	1	1	1	1	1	N/A	5
	FC	1	1 or 2	N/A	1	1	N/A	4 or 5
Production	iSCSI	1	1 or 2	2	1	1	N/A	6 or 7

Table 4) Clustered Hyper-V servers using live migration and CSVs configuration.

Environment	Protocol	Mgmt	VMs	iSCSI	Cluster	Migration	CSV	Total
Nonproduction	FC	1	1	N/A	1	1	1	5
	iSCSI	1	1	1	1	1	1	6
	FC	1	1 or 2	N/A	1	1	1	5 or 6
Production	iSCSI	1	1 or 2	2	1	1	1	7 or 8

Network Adapter Teaming

In the past, there has been much confusion about the support of (NIC) teaming with Microsoft Hyper-V. Microsoft explicitly states that it does not support use of NIC teaming with Microsoft Hyper-V. This means that Microsoft does not have a proprietary driver for all types of NICs installed within the Hyper-V server that supports NIC teaming. Only the manufacturers of the NICs, for example, Intel and Broadcom, have these drivers, and their driver software supports the use of NIC teaming.

Therefore, you can use NIC teaming in Windows Server 2008 R2 as long as the manufacturer of the NIC supports it. Microsoft does not contest enabling the Hyper-V role and then assigning a logical network interface that represents the teamed NICs to a virtual switch. The benefits of this are somewhat obvious: not only does the virtual switch have access to increased bandwidth, but it also has access to increased availability thanks to the redundancy of teaming multiple NICs together. However, some of the network features are not available on the logical NIC that is created as a result of teaming one or more physical NICs; this is discussed more under “Network Feature Support” later in this section.

Microsoft does not support the use of NIC teaming for networks used for iSCSI communications. Therefore, you might not use a logical network interface for iSCSI communications when using the Microsoft iSCSI Software Initiator to present LUNs to the Hyper-V parent partition. In addition, when a logical NIC interface is assigned to a virtual switch, the VMs that are connected to that specific virtual switch must not have the Microsoft iSCSI Software Initiator enabled for that virtual NIC.

Best Practice

For functional areas with more than one connection, such as the multiple network adapters used for VM communications, the connections should be spread across different network adapters, especially if multiple port network adapters are installed. This allows those functional areas to maintain connectivity to the network when configured properly so that, in the event of a port or adapter failure within the Hyper-V server, connectivity is not lost.

Since NIC teaming is not supported for iSCSI communications, NetApp recommends configuring multiple paths for storage connectivity to provide redundancy and additional bandwidth in some configurations using multipathing/MPIO. For more information on use of multiple paths for storage connectivity, see “Multipathing/MPIO” in section 5.1.

Designing for Bandwidth

Now that 10GbE is becoming more common in the data center, network design is changing. Because of the large leap in bandwidth, many are able to reduce the total number of NICs installed in servers and still support the network requirements for those servers. However, a Hyper-V server is much different from other servers, and therefore the addition of 10GbE NICs does not necessarily mean that a reduction in the total number of physical NICs will follow.

Most functional areas require a dedicated physical network adapter. Whether the dedicated network interface is a GbE or 10GbE NIC does not matter in most cases. In fact, the functional areas listed in Table 1 through Table 4 are not constrained by bandwidth in most environments: Hyper-V management (mgmt), cluster, and migration. Therefore, these functional areas benefit the least from increased bandwidth.

However, the following functional areas would benefit the most from increased bandwidth:

- iSCSI
- CSV
- VMs (in some cases)

The interface used for CSV communication would benefit most from the availability of increased bandwidth to support I/O redirection when it occurs.

Virtual Networks

You can create four different types of virtual networks within Hyper-V:

- Dedicated
- External
- Internal
- Private

For more information on the different types of virtual networks possible, see Configuring Virtual Networks on Microsoft TechNet. Also, for information on the setting within R2 to separate communications between the Hyper-V parent partition and the virtual networks, see Virtual Network Manager on Microsoft TechNet and New in Hyper-V Windows Server 2008 R2 Part 1: Dedicated Networks on John Howard's Microsoft TechNet blog.

Network Feature Support

With the release of Windows Server 2008 R2, support for several new networking features has been added, including jumbo frame support for GbE networks and TCP chimney support for 10GbE networks. These network technologies allow Hyper-V R2 to take advantage of network offload technologies, so instead of the Hyper-V CPUs processing network packets, these packets can now be handled by the offload technologies to help improve performance by reducing CPU use.

Windows Server 2008 R2 Hyper-V can fall victim to a few different performance bottlenecks within the network architecture and can be classified into two categories: receive path and transmit path. Table 5 lists possible reasoning for bottlenecks on the receive and transmit side.

Table 5) Hyper-V networking performance bottlenecks.

Receive and Transmit Path	Receive Path	Transmit Path
Data movement between parent and child partitions	Parsing packets to group based on MAC address	Task offload in software for VM-to-VM traffic

Receive and Transmit Path	Receive Path	Transmit Path
MAC address lookup and VLAN ID filtering		Additional copy for VM-to-VM traffic
Parent/child context switch overhead		

Large Send Offload (LSO) and Checksum Offload (CSO)

Large send offload (LSO and checksum offload (CSO) are supported by the virtual networks within Hyper-V. In addition, if the physical network adapters support it as well, the virtual traffic is offloaded to the physical network as necessary. Most network adapters support LSO and CSO these days, but check with your NIC manufacturer to be sure. If it is supported, it is often enabled by default.

Best Practice
Where LSO and CSO are supported, NetApp strongly recommends making sure that they are enabled (if not already enabled by default).

Jumbo Frames

Jumbo frames support was initially added win the introduction of Windows Server 2008, but additional improvements have been made to supportability with R2, and the NIC vendors have added a wider variety of NICs that support it. With Windows 2008 R2, enhancements converge to support up to six times the payload per packet, making a huge difference in overall throughput and reducing CPU utilization for large file transfers. In addition, jumbo frames are supported not only on the physical network adapters (as with Windows 2008 before R2), but also on the virtual networking, including switches and adapters. NetApp recommends use of jumbo frames with NICs configured to handle iSCSI communications, but they should be enabled only if there is end-to-end support across all hardware.

Use `ping -l 8000 -f -n 1 <target-ip>` to help determine if there is end-to-end support. Using an 8,000-byte datagram does not make it fragment at layer, and use of the -f option prevents false positives.

TCP Chimney

TCP chimney offload supports making connection offload capabilities available to the TCP/IP stack in the child partition. The virtual NIC in the child partition advertises the connection offload capabilities, and then the VM switch in the parent partition offloads the child partition TCP connections to the physical NIC. Applications with long-lived connections with large data transfers and applications with preposted buffers benefit the most from TCP chimney support.

Overall, the major benefits of TCP chimney offload support in a Microsoft Hyper-V virtual environment are as follows:

- Significant CPU utilization reduction on end-to-end workloads
- 10GbE connections can be fully utilized
- Avoids excessive chatter between child and parent partitions
  - Avoids the overhead of parent/child context switching
  - The connection state is fully maintained by the physical NIC
- Live migration support; connections are uploaded to the host stack during the live migration process

Best Practice

Where TCP chimney is supported by the NIC manufacturer, NetApp recommends making sure that it is enabled. (It is disabled by default.)

Virtual Machine Queue

Virtual machine queue (VMQ) helps significantly by classifying and grouping the received packets, parsing them in the hardware, and thereby improving performance by doing so. It also applies VLAN filtering in the hardware, dropping all packets with invalid VLAN IDs at the NIC, also using a switch on the NIC to do route lookup on transmits, and avoids a copy of NIC receive buffers to VM address space. All of the VMQ processing happens concurrently in multiple queues that are serviced by different processors.

NICs that support VMQ have embedded switches that allow receive queues to be paired with transmit queues, where each queue pair is a switch port. The switch requires no MAC address learning, and the switch inspects all transmit packets for destination MAC address + VLAN ID. If the packets pass the filter set on the receive queue, they DMA to that queue; otherwise they are sent on the wire. This is a huge advantage in VM-to-VM communication because it avoids route lookup in the software, avoids packet copies, and takes advantage of offload support in the hardware. Even with VM-to-physical communication, route lookup is still avoided, providing some benefit to performance.

Overall, VMQ improves network throughput by distributing network traffic for multiple VMs across multiple processors, while reducing processor utilization by offloading packet classification to the hardware and avoiding both network data copy and route lookup on transmit paths. VMQ is compatible with most other task offloads, and therefore it can coexist with large send offload and jumbo frames, but where TCP chimney is supported by the NIC, VMQ takes precedence. VMQ is secure and supports live migration in R2. By far, the best performance gains are seen with VMQ enabled on 10GbE network interfaces.

VMQ is supported only with specific adapters in Windows Server 2008 R2, and therefore VMQ is disabled by default. Before enabling it, check with Microsoft and your NIC vendor to make sure that your configuration supports VMQ.

Best Practice

Where VMQ is supported by the NIC manufacturer, NetApp recommends considering enabling it in your environment, especially if you have deployed Hyper-V with 10GbE.

Network Naming Standard

Network Naming Considerations

Because there are so many physical and virtual network adapters within a specific Hyper-V server, managing them is difficult. Most administrators establish an easy-to-understand naming standard for all network adapters, both physical and virtual. When deciding on a naming convention consider the following points:

- The name should identify whether the adapter is a physical or virtual network adapter.
- The naming convention should standardize the number of characters allowed, regardless of whether it is a physical or a virtual network adapter.
- For a physical network adapter, identify the physical network adapter hardware.
  - If the physical network adapter is located on the motherboard of the server, consider abbreviating the following: <M for motherboard><Double-Digit Port #>.
  - If the physical network adapter is an add-on to the server, consider abbreviating the following: <A for Add-on><PCIe/x Slot #><Single-Digit Port #>.

- For a physical network adapter connected to a physical network:
  - Use an abbreviated descriptor such as LAN for local area network, SAN for IP storage network, and so on.
  - If using VLAN tagging, use the VLAN ID of the network.
  - Create an abbreviation for the network subnet, using letters for class and three digits for the subnet identifier.
- For a physical network adapter connected to a virtual network/switch:
  - Use an abbreviated descriptor such as D for dedicated, E for external, I for internal, and P for private.
  - Use a two-digit code to differentiate the virtual network types, as you might often have multiple virtual networks of the same type.
- If it is a virtual or physical network adapter connected to an external or dedicated virtual network, identify the virtual network type.
  - Use an abbreviated descriptor such as D for dedicated, E for external, I for internal, and P for private.
  - Use a two-digit code to differentiate the virtual network types, as you might often have multiple virtual network of the same type.
- If it is a virtual or physical network adapter connected to an external or dedicated virtual network, then describe the network to which it is connected.
  - Use an abbreviated descriptor such as LAN for local area network, SAN for IP storage network, and so on.
  - If using VLAN tagging, use the VLAN ID of the network.
  - Create an abbreviation for the network subnet. First, use a single alpha character to identify the subnet class. Second, use two, three, or five number characters to identify the subnet, with three digits for the class octet and/or two digits for the subnet mask.

Network Naming Standard Examples

These network naming standard suggestions are used to establish a network naming convention. Using the simple naming convention works best for environments that will make use of 802.1Q VLAN trunking to minimize the number of physical network adapters required in the Hyper-V server and for most environments. An example of a complicated naming standard makes it clear that getting too complicated can actually be limiting in your environment and becomes unnecessary.

- Broadcom BCM5708C GbE, one of two ports located on the motherboard, connected to the management network on VLAN 15 which has a Class A subnet of 10.0.15.x/8.
  - Simple: PM01-MGT15 = <P for physical><M for motherboard><Double-Digit Port #>-<MGT for Management Network><VLAN ID>
  - Complicated: PM01-MGT15-C01008 = <P for physical><M for motherboard><Double-Digit Port #>-<MGT for Management Network><VLAN ID>-<Subnet Class><Subnet Identifier-Class Octet><Subnet Identifier-Subnet Mask>
- Intel® PRO/1000 PT Quad Port LP Server Adapter installed in PCI Add-on slot 2, two of four ports located on the card, connected to the production network on VLAN 10, which has a Class B subnet of 172.16.100.x/16, and is used by External Virtual Network #2.
  - Simple: PA22-E0210 = <P for physical><A for Add-on><PCIe/x Slot #><Single-Digit Port #>-<E for External Virtual Network><two-digit virtual network id><VLAN ID>
  - Complicated: PA22-E0210-B1628 = <P for physical><A for Add-on><PCIe/x Slot #><Single-Digit Port #>-<E for External Virtual Network><two-digit virtual network id><VLAN ID>-<Subnet Class><Subnet Identifier-Class Octet><Subnet Identifier-Subnet Mask>

- Internal Virtual Network (Adapter) #1, connecting VMs on the SQL Network on VLAN 18, which has a Class C subnet of 192.168.80.x/24
  - Simple: VI01-SQL18 = <V for virtual><I for internal><two-digit virtual network id>-<SQL for Production Network><VLAN ID>
  - Complicated: VI01-SQL18-C08024 = <V for virtual><I for internal><two-digit virtual network id>-<SQL for Production Network><VLAN ID>-<Subnet Class><Subnet Identifier-Class Octet><Subnet Identifier-Subnet Mask>

Establish an easy-to-understand naming standard for all network adapters, both physical and virtual. Minimizing the number of characters in the standard is advantageous. Not only does this assist in improving management of the many networks within a Hyper-V server, but for environments that make use of scripting, an easy-to-understand naming standard is particularly advantageous.

Keep track of the physical and virtual network adapter configurations and related network information in a document, spreadsheet, or database. This becomes extremely important in larger environments and those deploying Hyper-V servers that are part of a failover cluster to make sure that all virtual networks are named exactly the same between all Hyper-V cluster nodes.

Network Adapter Binding Order and Metric Values

Network Adapter Binding Order

Because most Hyper-V servers have a multitude of network adapters, both physical and virtual, often the network adapter binding order might be configured incorrectly by Windows. This requires that the administrator verify that the network adapter binding order is correct for each Hyper-V server. This is especially important for Hyper-V servers configured as part of a Windows failover cluster. Modifying the network adapter binding order can be accomplished through Network Connections, under Advanced > Advanced Settings, in the Connections field. Network Connections can be found under Control Panel > Network and Internet in Windows Server 2008.

For server core environment, the nvspbind tool can be used to modify network bindings from the command line. For more details, refer to http://archive.msdn.microsoft.com/nvspbind.

Best Practice

After all the physical network adapters have been installed, all the Hyper-V virtual networks have been created, and all the networks have been named according to any standard, you must modify the network adapter binding order appropriately.

The first adapter in the binding order should be the adapter used for managing the Hyper-V parent partition. The adapters for iSCSI, live migration, and CSV communications should follow next, with the private network used for cluster heartbeat and all adapters associated with virtual networks done last.

Network Adapter Metric Values

Changing network metric values is not necessarily required for Hyper-V servers that are not members of a Windows failover cluster, and for those that are, it is also not necessary to make network metric changes for nonproduction environments.

For production Hyper-V servers that are members of Windows failover clusters, changing network metric values should be considered on a case-by-case basis, primarily to make sure that the Hyper-V parent partition does not prefer any other network over the dedicated physical Hyper-V parent network adapter.

Modifying the network adapter metric values can be accomplished through Network Connections, under <NIC> Properties > IPv4 Properties > Advanced TCP/IP Settings, in the Automatic Metric field. A value of 9999 assigns the network adapter the highest user-assignable link cost.

Table 6) Recommended network binding order and metric values.

Functional Network Description	Binding Order Priority	Metric Value
Hyper-V management	1	100–199
iSCSI	2	200–299
Live migration	3	400–499
Cluster shared volumes	4	300–399
Other networks	5	1,000–4,999
Cluster heartbeat	6	500–599
Virtual machines (pNICs)	7	5,000–6,999
Virtual switches	8	7,000–9,999

3.2 Hyper-V Cluster Network Considerations

The best practices discussed earlier are very important but are elevated even more in importance, considering the deployment of Hyper-V servers as part of a failover cluster.

After the failover clustering has been enabled on all Hyper-V servers configured as part of the cluster, additional configuration is still recommended for optimal performance.

Preconfiguration Recommendations

Before failover clustering is enabled on any Hyper-V server, there are considerations that need to be made and specific steps to be taken with regard to the Hyper-V server's network configuration.

1. Ideally, all physical NICs, across all Hyper-V servers that are to be part of the new failover cluster, should be named exactly the same. If you are following the network naming standard discussed earlier in section 3.1, this is possible only if you are using the same server hardware and the same physical NICs, and you are installing the network PCI cards in the same slots across all servers in the cluster. When you are not using the same server hardware and physical NIC configuration, you might consider simplifying the network naming standard even further to allow all network connections to be named similarly. However, this is not a requirement, only a suggestion for consideration.
2. Make sure that the network connections considered as part of the same functional group, as initially defined in Table 6, are configured similarly. Minimally, each NIC should be able to communicate with other NICs in the same functional group and should be configured with similar IP settings (except for IP address). For the best possible outcome, the physical NICs should have not only the similar IP settings but also the same hardware settings within the adapter properties themselves.
3. Make sure that the network binding order is the same for each Hyper-V server to be configured as part of the Windows failover cluster.
4. Similarly, if you modify the network metric values for the physical NICs, then these should be the same across all physical NICs, across all Hyper-V servers to be configured as part of the Windows failover cluster.

NetApp recommends using the preceding steps as an opportunity to check over your network configuration, even if you have already followed these best practices as part of the initial configuration of your physical NICs. Having a consistent configuration across all Hyper-V nodes within the cluster allows for optimal performance and ease of management.

NetApp asks administrators to consider disabling all noncritical physical NICs before enabling failover clustering. Critical NICs include those used for Hyper-V management, IP storage communications (if present), and the cluster private network (commonly used for Windows failover cluster heartbeat

communications). After all noncritical NICs have been disabled, you may proceed with adding the failover cluster feature to the Hyper-V server.

Postconfiguration Recommendations

After the failover cluster feature has been added, use Failover Cluster Manager to begin configuring your cluster networks and proceed to the following section, “Naming Cluster Networks.”

If you have disabled all noncritical NICs before enabling the failover cluster feature, then the only networks present under Networks within Failover Cluster Manager will be the critical networks that were not disabled prior to configuring the Hyper-V server with failover clustering. Therefore, continue to follow the steps in this section before proceeding to the following section, “Naming Cluster Networks.”

Before making any further changes, begin by enabling the same adapters across all Hyper-V nodes within the failover cluster. However, enable only one functional network at a time, which means you may enable only one or two physical NICs per Hyper-V node at a time, repeating this for all Hyper-V nodes within the cluster. Continue the process to enable physical NICs that are part of the same logical network (usually defined by functional areas as laid out in Table 6), one cluster network at a time, for each Hyper-V node in the cluster. After all physical NICs across all Hyper-V nodes have been enabled for a specific cluster network, return to Failover Cluster Manager to make sure that they all are grouped within the same cluster network as previously planned.

Although you might choose to enable all the previously disabled NICs at the same time, if an issue arises in which the cluster networks are not defined as planned across all Hyper-V nodes within the cluster, then fixing the cluster networks is only further complicated by the rush.

Naming Cluster Networks

Use the same general recommendations for naming clusters as for naming networks (as discussed in section 3.1 under “Network Naming Standard”).

Best Practice

Establish an easy-to-understand naming standard for all cluster networks, one that contains as few characters as possible but describes the function of the network.

For example: In NetApp internal labs, we use `Mgmt` as a name for the cluster networks composed of the NICs used to manage the Hyper-V nodes, `Storage` for the cluster network used for iSCSI communications, and `Private` for the cluster network used for the failover cluster heartbeat.

Configuring Cluster Networks

After all of the cluster networks are properly defined for the entire Hyper-V cluster, we can now proceed to configure the settings for each cluster network. This can be done by viewing each cluster network within Failover Cluster Manager and then choosing from the pane on the right side of Failover Cluster Manager titled Actions, or by right-clicking a specific cluster network and selecting Properties.

Figure 2) Cluster network properties.

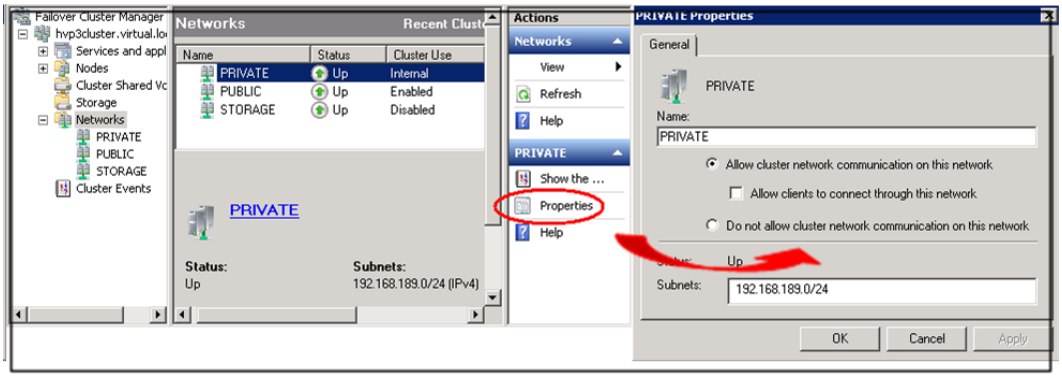

Use Table 7 to configure each cluster network according to its function.

The settings are abbreviated as follows:

- Allow cluster… = Allow cluster network communication on this network
- Allow clients… = Allow clients to connect through this network
- Do not allow cluster… = Do not allow cluster network communication on this network

Table 7) Recommended cluster network configuration settings.

Cluster Network Description	Cluster Network Setting
Hyper-V management	Allow cluster… /Allow clients…
iSCSI	Do not allow cluster…
Live migration	Do not allow cluster…
Cluster shared volumes	Allow cluster…
Other networks	Configure based on environment
Cluster heartbeat	Allow cluster…
Virtual machines (pNICs)	Do not allow cluster…
Virtual switches	Do not allow cluster…

Best Practice

NetApp provides these configuration settings for cluster networks as a reference but understands that each Hyper-V deployment differs from the others. This includes those who haven't configured each functional area listed in Table 7 to use a separate physical network. Therefore, use your best judgment when considering the configuration recommended.

Where the configuration in Table 7 matches your environment, NetApp recommends applying the configuration settings shown in Table 7. However, NetApp does not expect cluster network names to match the functional description provided in Table 7. Use your best judgment when matching your cluster networks to those described functionally in this table.

Cluster Network Metrics

Although CSVs are discussed in more detail later, we have discussed the recommendation from Microsoft to dedicate a minimum of one physical NIC for CSV communications.

Best Practice

NetApp supports Microsoft's recommendation to dedicate a minimum of one physical NIC for CSV communication in each Hyper-V node within the cluster.

Each cluster network within a Hyper-V cluster has two settings for network prioritization: Metric and AutoMetric. The Metric value is used to determine the priority for a specific cluster network. The AutoMetric setting is True or False, depending on whether the administrator configures a manual setting for the Metric value. For private cluster networks, the Metric value should be 1,000 to 10,000; for public cluster networks, the Metric values begin at 10,000.

The Hyper-V cluster uses the Windows PowerShell cmdlet `Get-ClusterNetwork` to prioritize the cluster networks and choose the cluster network with the lowest Metric value as the preferred network for CSV communication. Why is it preferred rather than dedicated? This network connection used for CSV communication is fault tolerant. Therefore, if the primary cluster network configured for CSV communications experiences issues or fails altogether, then the Hyper-V cluster detects this event and automatically moves CSV communications to the cluster network with the next lowest Metric value.

The same Windows PowerShell cmdlet used to identify the cluster network used for CSV communication (`Get-ClusterNetwork`) also allows us to change the settings of the cluster networks. Using the Windows PowerShell cmdlet is the only way to configure a preferred network for CSV communication; in R2, there is no option within any user interface to configure the cluster network Metric values. Table 8 provides the recommended cluster network values.

Table 8) Recommended cluster network AutoMetric and metric values.

Cluster Network Description	AutoMetric	Metric
Hyper-V management	True	Auto
iSCSI	True	Auto
Live migration	False	1,000–1,099
Cluster shared volumes	False	500–599
Other networks	True	Auto
Cluster heartbeat	False	1,500–1,999
Virtual machines (pNICs)	True	Auto
Virtual switches	True	Auto

Best Practice

NetApp provides these configuration settings for cluster networks as a reference, but we understand that each Hyper-V deployment differs from the others. This includes those who have not configured each functional area listed in Table 8 to use a separate physical network; therefore, use your best judgment when considering the configuration recommended in Table 8.

Where the configuration in Table 8 matches your environment, NetApp recommends applying the configuration settings shown. However, NetApp does not expect cluster network names to match the functional description provided in Table 8. Use your best judgment when matching your cluster networks to those described functionally in this table.

Table 8 assumes that each cluster network described in the left column is a separate physical network. If you have combined functions on the same network or are using VLANs to define separate logical networks versus physical networks, use your best judgment when setting Metric values for each cluster network configured within the Hyper-V cluster.

Table 8 also assumes that a Metric value between 500 and 999 assigned to the cluster network preferred for CSV communications is the lowest setting of all cluster networks present. It should not be possible for any other private cluster network to have a value lower than 1,000; therefore, you may leave the AutoMetric setting for all other cluster networks set to Auto as long as you confirm that the cluster network you prefer for CSV communications is configured with the lowest Metric value. If not, consider configuring additional cluster network metrics to make sure there is preference of cluster networks for CSV and live migration communications.

To manage the cluster network used for CSV communications, follow these steps:

1. Log in to any Hyper-V Server that is a member of the failover cluster.
2. Open Windows PowerShell Modules in Administrative Tools: Start > All Programs > Administrative Tools > Windows PowerShell Modules.
3. To enumerate the current cluster network settings, type the following command:

```
Get-ClusterNetwork | ft Name, Metric, AutoMetric
```

4. To set the metric value for a specific cluster network, you must know the name of the cluster network, taken from the information provided in the results of the previous command. After you know the name of the cluster network you want to configure for CSV communication, type the following commands:

```
$csv = Get-ClusterNetwork "[insert cluster network name]"
$csv.Metric = [insert recommended metric value from Table 8]
```

5. To validate the current cluster network settings, type the following command:

```
Get-ClusterNetwork | ft Name, Metric, AutoMetric
```

3.3 Storage Networking Considerations

When you design a network infrastructure (FC or IP), it should have no single point of failure. A highly available solution includes having two or more FC or IP network switches, two or more host bus adapters (HBAs) or NICs per Hyper-V server, and two or more target ports or NICs per storage controller. In addition, if you are using Fibre Channel, two fabrics are required to have a truly redundant architecture. For additional information on designing and deploying an FC/iSCSI solution, refer to the NetApp Fibre Channel and iSCSI Configuration Guide on the NetApp Support site for your version of Data ONTAP®.

Fibre Channel Storage Networking

NetApp clustered storage systems have an option known as controller failover mode (cfmode) that defines how Fibre Channel ports behave during failover in an active-active configuration. Selecting the

right cfmode is critical to making sure your LUNs are accessible and optimizing your storage system's performance in the event of a failover.

Best Practice

NetApp strongly recommends that the cfmode be set to single system image (SSI) because it provides LUNs accessibility across all storage ports.

To verify the current cfmode using the NetApp console, follow these steps:

1. Log in to the NetApp console using either SSH, telnet, or console connection.
2. Type the following into the prompt:

```
fcp show cfmode
```

3. If cfmode must be changed to SSI, type the following into the prompt:

```
priv set advanced
```

4. Type the following into the prompt:

```
fcp set cfmode <mode type>
```

Single-system image requires additional multipathing configuration on the Hyper-V server. For more information about the different cfmodes available and the impact of changing a cfmode, see the NetApp Block Access Management Guide for iSCSI and FC on the NetApp Support site for your version of Data ONTAP. For a complete list of supported Hyper-V configurations, see the NetApp Interoperability Matrix.

Best Practice

The storage controllers should have at least two FC HBA ports available to make redundant paths available between the NetApp storage system and the Hyper-V servers.

IP Storage Networking

NetApp Virtual Interfaces

A virtual network interface (VIF) is a mechanism that supports aggregation of network interfaces into one logical interface unit. Once created, a VIF is indistinguishable from a physical network interface. VIFs are used to provide fault tolerance of the network connection and in some cases higher throughput to the storage device.

Multimode VIFs are compliant with IEEE 802.3ad. In a multimode VIF, all of the physical connections in the VIF are simultaneously active and can carry traffic. This mode requires that all of the interfaces are connected to a switch that supports trunking or aggregation over multiple port connections. The switch must be configured to understand that all the port connections share a common MAC address and are part of a single logical interface. Figure 3 is an example of a multimode VIF; interfaces e0, e1, e2, and e3 are part of the MultiTrunk1 multimode VIF. All four interfaces in the MultiTrunk1 multimode VIF are active.

In a single-mode VIF, only one of the physical connections is active at a time. If the storage controller detects a fault in the active connection, then a standby connection is activated. No configuration is necessary on the switch to use a single-mode VIF, and the physical interfaces that make up the VIF do not have to connect to the same switch. IP load balancing is not supported on single-mode VIFs. Figure 4 is an example of a single-mode VIF; in the figure, e0 and e1 are part of the SingleTrunk1 single-mode VIF. If the active interface (e0) fails, the standby e1 interface takes over and maintains the connection to the switch.

Figure 3) Multimode VIF.

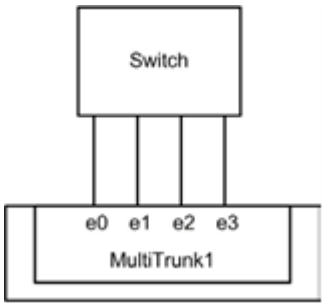

Figure 4) Single-mode VIF.

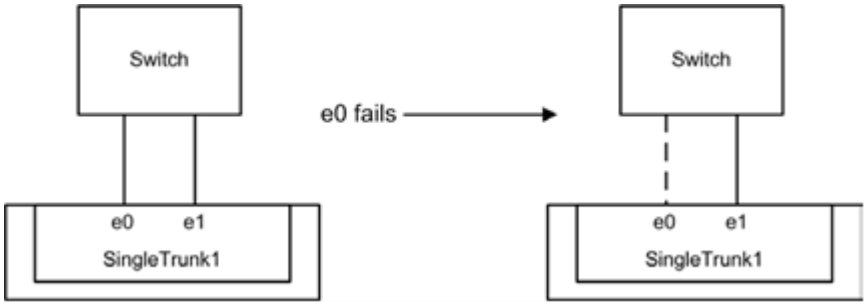

It is also possible to create second-level single or multimode VIFs, as shown in Figure 5. By using second-level VIFs, it is possible to take advantage of both the link aggregation features of a multimode VIF and the failover capability of a single-mode VIF. In this configuration, two multimode VIFs are created, each one to a different switch. A single-mode VIF is then created composed of the two multimode VIFs. In normal operation, traffic flows over only one of the multimode VIFs, but in the event of an interface or switch failure, the storage controller moves the network traffic to the other multimode VIF. For detailed information on the different types of VIFs, refer to the Data ONTAP Network and File Access Management Guide for your version of Data ONTAP.

Figure 5) Second-level VIF.

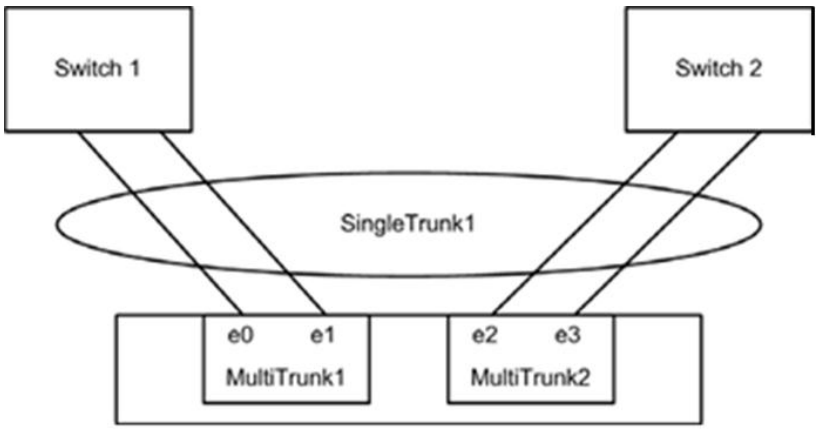

Best Practice

The storage controllers should have two or more target ports to be sure that there are redundant paths available between the NetApp storage system and the Hyper-V servers.

The use of LACP (dynamic multimode VIF) is also supported from Data ONTAP 7.2.1.

However, dynamic multimode VIFs have some special requirements:

- Dynamic multimode VIFs must be connected to a switch that supports LACP.
- They must be configured as first-level VIFs.
- They should be configured to use the IP-based load-balancing method.

iSCSI Traffic Security

NetApp storage controllers also allow the restriction of the iSCSI protocol to specific interfaces or VLAN tags. These simple configuration settings have an enormous effect on the security and availability of IP-based host disks.

Best Practice

For Hyper-V environments that will be deployed using iSCSI, NetApp strongly recommends that separate physical networks be implemented for handling the general IP communication between servers and virtual machines and for handling all storage-related IP communication. At minimum, the traffic should be virtually segregated using 802.1Q VLAN tagging.

Masking and Zoning

When storage is provisioned for access using FCP or iSCSI, the storage must be masked so that the appropriate Hyper-V parent and child partitions can connect to it. With a NetApp storage system, LUN masking is handled by the creation of initiator groups (also known as igroups).

If NetApp SnapDrive® for Windows is installed within the Hyper-V server or a guest OS, then it can be used to configure igroups through the Manage IGroup wizard under Disks > Actions. However, if you are not using NetApp SnapDrive, then a more manual process must be followed in which you first obtain the iSCSI qualified names (IQNs) or worldwide port names (WWPNs) for the configured storage connectivity. This is discussed in detail later.

Best Practices

NetApp recommends creating an igroup for each Hyper-V server, each Windows failover cluster (a combination of multiple Hyper-V servers), or each guest OS (for the option of direct LUN access by guest OS using the Microsoft iSCSI Software Initiator), depending on requirements.

Note: If a Hyper-V server or cluster uses both Fibre Channel and iSCSI protocols, separate igroups must be created for Fibre Channel and iSCSI.

NetApp also recommends including the name of the Hyper-V server, Windows failover cluster, or guest OS and the protocol type (for example, `DC1_FCP` and `DC1_iSCSI`) within the naming convention for the igroup. This naming convention and method simplify the management of igroups by reducing the total number created. It also means that all Hyper-V servers in the Windows failover cluster see each LUN at the same ID. Each initiator group includes all of the FCP WWPNs or IQNs of the Hyper-V servers in the cluster.

NetApp recommends using NetApp SnapDrive to manage igroups for the configured storage.

Obtaining the IQN (iSCSI)

Obtain the IQN from the Hyper-V server or guest OS by opening a command prompt and running the following command: **`iscsicli`**. For more information on obtaining the IQN for all configured iSCSI initiators, see "Configure Windows Server 2008 Initiator Groups on NetApp Storage" in TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide and follow the instructions.

Obtaining the WWPN (FC)

Obtain the WWPN for the FC ports using the NetApp Windows Host Utility (discussed later in this TR) or the HBA vendor host utility (for example, Qlogic SAN Surfer) on the Hyper-V server. For more information on obtaining the WWPN of all Fibre Channel HBAs installed, see “Fibre Channel Zoning Configuration” in TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide and follow the instructions.

4 Storage Configuration

In a Hyper-V environment, the availability and performance of the shared storage infrastructure are more critical than those of the individual servers running the virtualized server environment. It is therefore vital that the required level of availability and performance be factored in when selecting and designing the storage solution for the virtualized server environment. NetApp offers a comprehensive set of software and hardware solutions to address the most stringent requirements for availability and performance of large, scalable Hyper-V environments.

4.1 Active-Active NetApp Controllers

When designing a shared storage environment for Microsoft Hyper-V, it is essential that the solution be highly available. NetApp uses an active-active controller design to make sure of data availability in business-critical environments, such as Microsoft Hyper-V virtual environments. Active-active controllers provide simple, automatic, and transparent failover to deliver enterprise-class availability. Providing the highest level of availability of the shared storage is critical as all servers depend on it. For more details, see High-Availability System Configuration and TR-3450: Active-Active Controller Overview and Best Practices Guidelines.

Best Practice
Use active-active storage configurations to improve overall system availability by eliminating a single point of failure (SPOF).

4.2 Multipath HA

Multipath HA storage configuration further enhances the resiliency and performance of active- active controller configurations. Although cluster failover software provides high availability by providing fault tolerance in the event of a controller failure, storage-triggered events often result in unneeded failovers or prevent successful takeovers. Multipath HA storage enhances storage resiliency by reducing unnecessary takeover by a partner node due to a storage fault, improving overall system availability and promoting higher performance consistency. Multipath HA provides added protection against various storage faults, including HBA or port failure, controller-to-shelf cable failure, shelf module failure, dual intershelf cable failure, and secondary path failure. Multipath HA helps provide consistent performance in active-active configurations by providing larger aggregate storage loop bandwidth. For details, see TR-3437: Storage Subsystem Resiliency Guide.

4.3 RAID Data Protection

A challenge of any server consolidation effort (for example, server virtualization) is increased risk if the consolidated platform fails. As physical servers are converted to virtual machines (called child partitions in Hyper-V terminology), and multiple virtual machines are consolidated onto a single physical server (called the parent partition in Hyper-V terminology), the impact of a failure to the storage platform could be catastrophic.

When focusing on storage availability, consider that many levels of redundancy are available for deployments, including purchasing physical servers with multiple storage interconnects or HBAs,

deploying redundant storage networking and network paths, and leveraging storage arrays with redundant controllers. A deployed storage design that meets all of these criteria can help to mitigate single points of failure. The reality is that data protection requirements in a Hyper-V virtual infrastructure are greater than those in a traditional physical server infrastructure, considering that a single server failure could result in the failure of multiple applications. Data protection is a paramount feature of shared storage devices.

NetApp RAID-DP® is an advanced RAID technology that provides the default RAID level on all storage systems. RAID-DP protects against the simultaneous loss of two drives in a single RAID group. It is very economical to deploy; the overhead with default RAID groups is a mere 12.5%. This level of resiliency and storage efficiency makes data residing on RAID-DP safer than data residing on RAID 5 and makes RAID-DP more cost effective than RAID 10.

Best Practice

Use RAID-DP, the NetApp high-performance implementation of RAID 6, for better data protection on all RAID groups that will store virtual disks for the Hyper-V virtual machines.

4.4 Remote LAN Management (RLM) Card

The RLM card provides secure out-of-band access to the storage controllers that can be used regardless of the state of the controllers. The RLM offers a number of remote management capabilities for NetApp controllers, including remote access, monitoring, troubleshooting, logging, and alerting features. The RLM also extends AutoSupport™ capabilities of the NetApp controllers by sending alerts or notification using an AutoSupport message when the storage system goes down, regardless of whether the controller can send AutoSupport messages. These AutoSupport messages also provide proactive alerts to NetApp to help provide faster service. For more details, visit Remote LAN Management (RLM) on the NetApp Support site.

Best Practice

Use the latest storage controller, shelf, and Data ONTAP general deployment release available on the NetApp Support site. As a minimum, NetApp recommends using Data ONTAP release 7.3 or later with Hyper-V virtual environments.

5 Storage Provisioning

5.1 NetApp Storage Software and Tools

NetApp Windows Host Utilities Kit

Windows Host Utilities is NetApp software that modifies system settings so that the Hyper-V parent or child OS operates with the highest reliability possible when connected to NetApp SAN storage.

The utility checks for the appropriate operating system patches, verifies and updates registry keys to set proper timeout values, and sets HBA parameters. The updates to registry entries and HBA parameters enable the host to correctly handle storage system failover and reboot, as well as to specify LUN types claimed by Windows multipathing components. In addition, the diagnostic programs (data collection scripts) provide information about the server, storage system configuration, and Fibre Channel switches. The outputs of these scripts are useful for NetApp Customer Support to troubleshoot your configuration.

NetApp Windows Host Utilities 5.x is supported on both Hyper-V server and child operating systems. These child operating systems include Windows Server 2008/R2, Windows Server 2003, Windows XP,

and Windows Vista. It supports Fibre Channel, iSCSI, and mixed FC and iSCSI connections. It also supports iSCSI HBAs.

Best Practice

NetApp strongly recommends installing the Windows Host Utilities Kit on all Hyper-V servers.

- For Windows Server 2008, install Windows Host Utilities Kit version 5.1 or higher.
- For Windows Server 2008 R2, install Windows Host Utilities Kit version 5.2 or higher.

For details on downloading, installing, and configuring NetApp Windows Host Utilities 5.x, read the Windows Host Utilities 5.0 Installation and Setup Guide, available on the NetApp Support site. In addition, for more information on the installation of NetApp Windows Host Utilities Kit 5.x, see "NetApp Windows Host Utilities Kit" in TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide and follow the instructions.

Multipathing/MPIO

The generic term "multipathing" is used to describe the ability of a system to use more than one read-write path to a storage device. Multipathing provides a high-availability solution that provides fault tolerance against a single point of failure in hardware components. Multipathing can also provide load balancing of I/O traffic, thereby improving system and application performance.

Microsoft Multipath I/O (MPIO)

Microsoft MPIO is a feature that provides support for using multiple data paths to a storage device. It increases availability by providing multiple paths (path failover) from a server or cluster to a storage subsystem. Windows Server 2008 R2 natively includes support for Microsoft MPIO.

Microsoft MPIO is supported by both the Hyper-V server parent partition (independent of the child OS) and by a supported Windows Server guest OS within a child partition, which does not include Windows XP or Vista.

Microsoft MPIO is protocol independent and is supported with iSCSI, Fibre Channel, and SAS. It is an extensible solution to enable storage partners such as NetApp to create device-specific modules (DSMs) to provide optimized solutions. The DSM, as its name suggests, is a module used to determine how MPIO should behave for specific devices (for example, NetApp LUN). Microsoft provides a generic DSM natively in Windows Server 2008 R2. Several storage path load-balancing policies are supported by the Microsoft DSM:

- Failover, in which no load balancing is performed
- Round-robin load balancing, in which the DSM uses all available paths for I/O in a balanced, round-robin fashion
- Round robin with a subset of paths load balancing, in which the application specifies a set of paths to be used in round-robin fashion and a set of standby paths
- Dynamic least queue depth load balancing
- Weighted-path load balancing

Regardless of the chosen setting, the Microsoft DSM remembers load balance settings across reboots.

NetApp has also developed a DSM called Data ONTAP DSM that helps MPIO interact with NetApp storage systems.

Data ONTAP DSM for Windows MPIO

The Data ONTAP DSM for Windows MPIO enables NetApp storage systems to integrate with Microsoft MPIO on Windows Server 2008 Hyper-V servers and thus provide high availability for applications. The

Data ONTAP DSM identifies its list of supported devices, and any time it sees a NetApp LUN, it claims that device (NetApp LUN). It determines all the paths pointing to the same LUN so that MPIO can group them into the virtual disk that Windows Server 2008 Hyper-V server will mount.

The Data ONTAP DSM is also responsible for communicating with MPIO about how to route I/O, especially important in the event of a failover. There can be multiple active paths and multiple passive paths. If all of the active paths fail, the DSM automatically switches to the passive paths, maintaining the host's access to its storage. For Windows Server 2008 Hyper-V servers, NetApp recommends Data ONTAP DSM 3.2R1 or higher. The installation program installs the DSM and any required Windows MPIO components. The program also sets Windows registry values and HBA parameters. It can coexist with the Microsoft MS DSM.

The Data ONTAP DSM is supported for both the Hyper-V server and the child OS, except for Windows XP and Vista because Microsoft MPIO is not supported for Windows XP and Windows Vista.

The Data ONTAP DSM can manage both FCP and iSCSI paths, including mixed FCP and iSCSI paths to the same virtual disk (LUN), and can apply load balancing as well as failover schemes to the LUN. The Data ONTAP DSM also supports multiple iSCSI paths (that is, two iSCSI HBAs, two NICS, one iSCSI HBA, and one NIC). At the time of publication, LUNs over iSCSI software initiator are supported at the child OS level.

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

- For Windows Server 2008, install Data ONTAP DSM version 3.2R1 or higher.
- For Windows Server 2008 R2, install Data ONTAP DSM version 3.3.1 or higher.

For the currently supported multipathing software versions and related requirements, see the NetApp Interoperability Matrix.

Data ONTAP Load-Balancing Policies

The Data ONTAP DSM for Windows MPIO supports a number of load-balancing policies for Windows Server 2008 Hyper-V server. These include the following:

- **Least queue depth.** An active-active policy that tracks the outstanding I/O queues across all paths and submits the current I/O to the path with the least queue depth. This policy provides the maximum performance/throughput and is the default policy.
- **Least weighted path.** An active-passive policy that allows the user to assign weights to each I/O path; the path with the least weight is used. Path weights persist across Windows Server 2008 reboots. A new path is chosen only when the active path fails.
- **Failover.** An active-passive policy in which all I/O is submitted to one path and is submitted to a different path only in case the primary path fails.
- **Autoassigned.** Similar to failover, with the exception that the administrator cannot change the active path and the DSM chooses which path should be used for all I/O.
- **Round robin**. An active-active policy in which I/O is sequentially submitted to all of the available paths.

Best Practice
The least queue depth policy is the default and the NetApp best practice recommendation for Hyper-V deployments with MPIO configured.

Data ONTAP DSM Coexistence with Other DSMs (iSCSI Only)

The Microsoft iSCSI software initiator also includes a DSM that can manage the iSCSI paths that use the software initiator. The two DSMs (Microsoft DSM and Data ONTAP DSM) can coexist, as long as both

versions are listed on the appropriate support matrixes. NetApp recommends that you install the Microsoft iSCSI DSM even if you claim iSCSI LUNs with the Data ONTAP DSM. When both DSMs are installed, the Data ONTAP DSM has priority in claiming iSCSI LUNs on NetApp.

If you install the Microsoft iSCSI software initiator on the same Hyper-V server as the Data ONTAP DSM, do the following:

- Install the Windows Host Utilities for native OS. For the recommended version, see section 5.1, "NetApp Storage Software and Tools."
- Check the support matrix for the latest information.
- If you want to use the Microsoft iSCSI DSM, be sure to install the Microsoft iSCSI DSM by selecting the Microsoft MPIO multipathing support for iSCSI checkbox when installing the initiator. Choose to claim only FCP LUNs when installing the Data ONTAP DSM.

For details on installing and configuring Data ONTAP DSM, read the Data ONTAP DSM for Windows MPIO Installation and Administration Guide, available on the NetApp Support site.

iSCSI Multiple Connections per Session

The iSCSI specification provides for an optional feature referred to as multiconnection sessions (MCS). iSCSI defines this concept as an iSCSI initiator establishing a session with an iSCSI target. MCS allows for this single iSCSI session to use multiple TCP connections. The use of MCS and MPIO for concurrently connecting to the same LUN is not supported.

NetApp SnapDrive for Windows

SnapDrive is NetApp software that helps a system administrator provision and manage storage directly from a server. SnapDrive also gives flexibility to application or system administrators by enabling them to define their data protection and business continuance policies, and more importantly, it allows administrators to resize the storage on the fly without any disruption of application service. SnapDrive simplifies storage and data management by using the host operating system and NetApp technologies, by hiding the complexity of steps that must be executed on both the storage system and the host system, and by removing the dependency on the storage administrator.

Key SnapDrive for Windows functionality includes SAN storage provisioning on the host, consistent data Snapshot copies, and rapid application data recovery from Snapshot copies. SnapDrive complements the native file system and volume manager technology, and it integrates seamlessly with the clustering technology supported by the host operating system to provide high availability of the service to its users. SnapDrive manages LUNs on a storage system, making these LUNs available as local disks on Windows hosts. This allows Windows hosts to interact with the LUNs just as if they belonged to a directly attached redundant array of independent disks (RAID).

SnapDrive provides the following additional features:

- Online storage configuration, LUN expansion, and streamlined management
- Connection of up to 128 LUNs
- Integration of Data ONTAP Snapshot technology, which creates point-in-time images of data stored on LUNs
- Integration with SnapMirror® software to facilitate disaster recovery from either asynchronously or synchronously mirrored destination volumes
- Enabling SnapVault® updates of qtrees to a SnapVault destination
- Management of SnapDrive on multiple hosts
- Simplified storage provisioning in Microsoft failover clustering configurations
- iSCSI session management

Best Practice

NetApp highly recommends the installation of NetApp SnapDrive on all Hyper-V and SCVMM servers to enable maximum functionality and support of key features.

For Microsoft Windows Server 2008 R2 installations in which the Hyper-V role is enabled and for Microsoft Hyper-V Server 2008 R2 to support:

- Existing features (no new R2 features), install NetApp SnapDrive for Windows version 6.1P2 or higher.
- New features (all new R2 features), install NetApp SnapDrive for Windows version 6.3P2

NetApp SnapDrive for Windows version 6.0 or higher is also supported for installation within the child OS on which Microsoft Windows Server 2003, Microsoft Windows Server 2008, and Microsoft Windows Server 2008 R2 are installed.

For detailed instructions on installation and administration (provisioning and managing LUNs, data protection, and so on), refer to the SnapDrive for Windows Installation and Administration Guide. In addition, for more information on the installation of NetApp Windows Host Utilities Kit 5.x, see the "NetApp SnapDrive for Windows" section of the TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide and follow the instructions within the sections "Download" and "Installation."

For the latest supported software versions and related requirements, see the NetApp Interoperability Matrix. Refer to the support matrix for further details on exact configurations supported.

5.2 NetApp Storage Provisioning

Aggregates

An aggregate is the NetApp virtualization layer, which abstracts physical disks on a storage device from logical datasets that are referred to as flexible volumes. Aggregates are the means by which the total IOPS available from all of the individual physical disks are pooled as a resource. The aggregate concept is well suited to meet users' differing security, backup, performance, and data-sharing needs, as well as the most unpredictable and mixed workloads. Aggregates can be 32-bit or 64-bit.

NetApp recommends that, whenever possible, a separate small aggregate with RAID-DP should be used for hosting the root volume. This aggregate stores the files required for running and providing GUI management tools for the NetApp storage system. The remaining storage should be placed into a small number of large aggregates; this provides optimal performance because of the ability of a large number of physical spindles to service I/O requests. On smaller arrays, it might not be practical to have more than a single aggregate because of the restricted number of disk drives on the system. In these cases, it is acceptable to have only a single aggregate.

64-Bit Aggregates

64-bit aggregates are supported from Data ONTAP 8.0.7 onward. In this type of aggregate, the aggregate size that can be configured is more than 16TB, which is the limit for 32-bit aggregates. 64-bit aggregates range from 40TB to 100TB, depending on the model of the storage system. 64-bit aggregates enable you to add more data disks to an aggregate than 32-bit aggregates. When you create a FlexVol volume inside a 64-bit aggregate, the FlexVol volume is striped across all the disks in the aggregate. Therefore, the FlexVol volume has the capability to use more spindles. This means that for workloads and storage system configurations that were constrained by the number of available disk drive spindles, 64-bit aggregates can provide better performance if you can add enough drives that the workload or storage system is no longer disk spindle constrained in delivering performance.

However, when a memory-intensive or highly random workload is running on 64-bit aggregates, it might have a slight performance impact. For more details on use cases and recommendations for 64-bit aggregates, refer to TR-3786: A Thorough Introduction to 64-bit Aggregates.

Best Practice

When you configure an aggregate, NetApp recommends that you use the following settings:

- **Double Parity**. Select this option to benefit from RAID-DP, which is the preferred RAID level for an aggregate.
- **RAID Group Size**. NetApp recommends selecting the default, which is 16 in most cases.
- **Disk Selection**. Automatic is selected by default and is the NetApp recommendation.
- **Disk Size**. By default, Any Size is selected. However, NetApp recommends selecting disks of the same size when creating an aggregate.
- **Number of Disks**. NetApp requires that at least three disks be assigned in order to provision a new aggregate. NetApp recommends creating the largest aggregate possible in order to benefit from the increased I/O capacity of all the spindles in the aggregate.

For details on configuring an aggregate, see the "Configure an Aggregate" section of TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide and follow the instructions.

Flexible Volumes

Flexible volumes contain LUNs that are accessed by Hyper-V servers over FC or iSCSI. They are virtual volumes that you can manage and move independently from physical storage, and they can be created and resized as your application needs change.

Best Practice

For configuring a FlexVol volume, NetApp recommends the following settings:

- **Volume Type Selection**. Select Flexible to create a NetApp FlexVol volume, which has many advantages over traditional volumes in a virtual environment.
- **Volume Name**. NetApp suggests using a combination of the Hyper-V hostname/cluster name and physical disk, the Hyper-V hostname/cluster name, physical disk, and replication policy: for example, `HostA_Disk2` and `HostB_Disk1_4hmirror`.
- **Language**. NetApp recommends accepting the default value, except where you have a good reason for changing it.
- **UTF-8**. NetApp recommends accepting the default value, except where you have a good reason for changing it.
- **Space Guarantee**. Choose the option desired. To achieve volume-level thin provisioning, select None. For more information, see the "Storage Thin Provisioning" section of this document.
- **Snapshot Reserve**. NetApp recommends that all volumes be configured with 0% and that the default Snapshot schedule be disabled. For details on disabling the default Snapshot schedule, see the "Create a NetApp Flexible Volume…" section of TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide and follow the instructions in step 7.

For details on configuring a NetApp flexible volume or FlexVol volume, see the "Configure a Flexible Volume" section of the TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide.

LUNs

LUNs are units of storage provisioned from a NetApp storage system directly to the Hyper-V servers. The Hyper-V server can access the LUNs using FC or iSCSI protocol as physical disks. The physical disks can serve the following purposes in a Microsoft virtualization environment:

- **Hyper-V parent partition**. In a boot-from-SAN deployment, Windows Server 2008 R2 can be installed on a LUN, provisioned, zoned (single path only during installation, return to MPIO after Hyper-V is installed), and connected to the physical Hyper-V server.
- **Windows failover clustering quorum or witness disk**. To use live migration and high availability for Hyper-V VMs, Windows failover clustering (WFC) must be configured. To complete the configuration of Windows failover clustering, a witness or quorum disk must be configured. All Hyper-V servers in the Windows failover cluster (WFC) must be able to access this disk; thus, provisioning a LUN from shared storage such as a NetApp storage array meets the requirements.
- **VHD storage**. A physical disk presented to the Hyper-V server can be used to store the virtual hard disk (VHD) files that a virtual machine uses as virtual disks for storage.
- **Virtual machine pass-through disk**. A physical disk presented to the Hyper-V parent partition in an offline state can be assigned to a virtual machine for use as a virtual disk for storage. Presenting the physical disk to the Hyper-V server in an offline state allows the child OS to use the disk while bypassing the Hyper-V parent partition file system, and it prevents the Hyper-V parent partition and child partition from accessing the physical disk at the same time.
- **Child OS storage**. A physical disk can be presented to the Hyper-V child partition directly, completely bypassing the Hyper-V parent partition storage architecture and file system. The child OS would use the Microsoft iSCSI Software Initiator to directly connect LUNs, using the iSCSI protocol and one or more of the virtual machine's virtual NICs. The child OS can then manage and directly access the physical disk for storing data the same way the Hyper-V parent partition would.

Process to Create a LUN

The procedure for creating LUNs is the same regardless of the protocol you use to access the LUN. However, you must create an initiator group (igroup) of the correct type for the protocol. LUNs can be created by using the Data ONTAP command line interface (CLI), by using the FilerView® interface, or by using NetApp SnapDrive where installed.

Best Practice

If NetApp SnapDrive for Windows is installed, NetApp strongly recommends that you use SnapDrive to configure initiator groups. Refer to the documentation for your version of SnapDrive for specific instructions. When configuring a LUN, NetApp recommends the following settings:

- **LUN Type.** Be sure to select the correct LUN type for your environment: `Dedicated` is appropriate for individual servers, and `Shared` is appropriate for clustered server nodes.
- **LUN Protocol Type.** For details on selecting the appropriate LUN type for your installation, see "Selecting the Correct LUN Protocol Type" following this section.
- **Space Reserved.** NetApp strongly recommends unselecting this setting, so that the LUN is not space reserved. For more information, see section 6.1, "Storage Thin Provisioning," in this report.

For details on configuring a new LUN, see the "Disk Provisioning on Windows 2008 Servers" section of TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide.

Best Practice

If SnapVault integration is planned, the LUN should always be placed in a qtree.

Selecting the Correct LUN Protocol Type

When creating the NetApp LUN using FilerView or the CLI, the parameter LUN type plays a very critical role. This parameter determines the on-disk layout of the LUN. It is important to specify the correct LUN type to make sure that the LUN is properly aligned with the file system on it. The underlying reason is that optimal performance with the storage system requires that I/O be aligned to a 4,096-byte boundary. Unaligned I/O might cause an increase in per-operation latency. Unaligned I/O requires the storage system to read from or write to more blocks than necessary to perform logical I/O. This issue is not unique to NetApp storage. Any storage vendor or host platform might exhibit this problem.

The LUN type you specify depends on the OS, OS version, disk type, and Data ONTAP version. If the incorrect LUN type is selected when the LUN is initially provisioned, a new LUN (with the correct LUN type) must be created, and the data must be transferred from the existing misaligned LUN. For complete information on LUN types for different OSs, see the Block Access Management Guide for your version of Data ONTAP.

Review Table 9 through Table 11 to determine the LUN types to use when configuring LUNs for use with Windows Server 2008, for installations with and without the Hyper-V role enabled.

Table 9) LUN types for use with Data ONTAP 7.3.1 and higher.

Data ONTAP 7.3.1 or Higher	Windows Server 2008 Physical Server w/o Hyper-V	Windows Server 2008 Hyper-V Server Physical Disk	Windows Server 2008 Hyper-V Server Pass-Through Disk
With SnapDrive installed	windows_gpt	hyper_v	LUN type of child OS
Without SnapDrive installed	windows_20008	windows_2008	LUN type of child OS

Table 10) LUN types for use with Data ONTAP 7.2.5 through 7.3.0.

Data ONTAP 7.2.5 Through 7.3.0	Windows Server 2008 Physical Server w/o Hyper-V	Windows Server 2008 Hyper-V Server Physical Disk	Windows Server 2008 Hyper-V Server Pass-Through Disk
With SnapDrive installed	windows_gpt	windows_gpt	LUN type of child OS
Without SnapDrive installed	windows_20008	windows_2008	LUN type of child OS

Table 11) LUN types for use with Data ONTAP 7.2.4 and earlier.

Data ONTAP 7.2.4 or Earlier	Windows Server 2008 Physical Server w/o Hyper-V	Windows Server 2008 Hyper-V Server Physical Disk	Windows Server 2008 Hyper-V Server Pass-Through Disk
With SnapDrive installed	Linux	Linux	LUN type of child OS
Without SnapDrive installed	Linux	Linux	LUN type of child OS

For Data ONTAP version 7.3 and earlier, the LUN type `windows_2008` is available only through the Data ONTAP CLI. Therefore, the LUNs for Hyper-V parent partition and Windows Server 2008 child VMs must be created through the LUN setup command on the Data ONTAP CLI.

For LUNs directly mapped to the child OS using the iSCSI software initiator running in the VM, the LUN type of the intended child OS should be selected when creating the LUN. In addition, for LUNs assigned directly to the child OS using a pass-through disk configured for the Hyper-V server, the LUN type of the intended child OS should be selected when creating the LUN.

NetApp requires that correct LUN protocol type be configured for all LUNs provisioned to the Hyper-V server. Failure to do so can result in misaligned VM file systems and mild to severe performance degradation.

If you have selected the incorrect LUN protocol type and then formatted the disk, the disk is not aligned to the NetApp storage array. You must create a new LUN, selecting the correct LUN protocol type, and then format the LUN. You must then move the VMs and VHDs off the LUN with the incorrect LUN protocol type. Using quick storage migration can help here, but downtime for the VM is required, with or without quick storage migration, in order to move the VMs and VHDs from the old LUN to the new LUN.

Do not just move the VM configuration files between LUNs. A couple of options for migrating the VMs between LUNs are:

1. You must first export the VMs before moving them. After the files have been moved, you can delete the old VM and import the VM again using the files from the previous process to export the VM.
2. Delete the VM and recreate it on the new LUN, then select the appropriate VHD, now located on the new LUN, as part of the process to create the new VM.

5.3 Microsoft Hyper-V Server Storage Provisioning

Microsoft Hyper-V supports the use of direct-attached storage, or DAS (SAS, SATA), and shared storage (through Fibre Channel or iSCSI only, no support for NAS at this time). Even with Hyper-V support for DAS, it quickly becomes clear when scaling your Microsoft Hyper-V environment beyond a few servers that the advantages of using shared storage far outweigh the lower acquisition costs of DAS. At minimum, implementing quick/live migration or high availability through Hyper-V integration with Windows failover clustering requires the presence of shared storage.

Traditional Volume

A traditional volume, for the purposes of this paper, is defined as a physical disk connected to the Hyper-V parent partition, and dedicated to that Hyper-V parent partition. A traditional volume might be created from DAS or shared storage, such as a LUN provisioned from the NetApp storage array.

Standard Cluster Volume

A standard cluster volume, for the purposes of this paper, is defined as a physical disk connected to the Hyper-V parent partition and connected to multiple Hyper-V server nodes configured as part of a single Windows failover cluster. However, only one Hyper-V server node in the cluster has access to read and write to the standard cluster volume, so that it might host a virtual machine (VM) and access the virtual hard disks or VHDs on the standard cluster volume. A standard cluster volume may be created only from shared storage, such as a LUN provisioned from the NetApp storage array, to which all Hyper-V server nodes in the failover cluster are connected.

Virtual Machine Limitations

When deploying Hyper-V with standard cluster volumes, Microsoft recommends a one-to-one mapping of virtual machines to SAN LUNs. Many refer to this as the Microsoft “one VM per LUN” best practice recommendation. The reasoning behind this is simple: if multiple VMs are present on a single standard cluster volume, the administrator cannot quickly migrate these VMs individually between Hyper-V server nodes in the cluster; instead, they are migrated as a group.

This is the result of the standard cluster volume’s being the root of every VM cluster resource; all VMs and VHDs present on the volume would need to be taken offline momentarily to migrate even one of the VMs (or all) between Hyper-V server nodes. In addition, Hyper-V is unable to migrate the VMs concurrently. Instead, they are migrated individually, and each VM must be placed into a saved state and subsequently returned to an operational state as part of the quick migration process.

Because downtime associated with an attempt to migrate a single VM increases with the presence of additional VMs on the same standard cluster volume, Microsoft has recommended a one-to-one mapping of virtual machines to SAN LUNs configured as standard cluster volumes. This remains true when deploying with Windows Server 2008 Hyper-V, as well as when deploying standard cluster volumes with Windows Server 2008 R2 Hyper-V.

Windows Server 2008 Hyper-V

In Windows Server 2008 Hyper-V, only quick migration is possible when migrating VMs and their resources (such as VHDs) between Hyper-V server nodes because only one Hyper-V server node is allowed to host the VM and access the VHDs on the standard cluster volumes. This means that downtime for the VM should be expected and taken into account when migrating VMs around the cluster.

Windows Server 2008 R2 Hyper-V

In Windows Server 2008 R2 Hyper-V, although only one Hyper-V server node is allowed to host the VM and access the VHDs on the standard cluster volumes, quick migration and live migration are both supported methods for migration of VMs in a pure Windows Server 2008 R2 Hyper-V cluster.

Live migration is a new feature with Windows Server R2, which allows the migration of a VM between Hyper-V server nodes within a Windows failover cluster without any interruption in service. Users connected to the VM that is being live migrated might notice a slight decrease in performance for a few moments, but ultimately the user's connection to the VM's applications or data should not be interrupted, and they should remain unaware that the VM was migrated from one physical Hyper-V server node to another. Live migration is possible with standard cluster volumes; a cluster shared volume is not a requirement to support live migration in a Windows Server 2008 R2 Hyper-V cluster. Microsoft still recommends a one-to-one mapping of virtual machines to standard cluster volumes.

Cluster Shared Volume

A cluster shared volume (CSV), for the purposes of this paper, is defined as a physical disk connected to the Hyper-V parent partition and shared between multiple Hyper-V server nodes configured as part of a single Windows failover cluster. This is a feature of failover clustering, available only with Windows Server 2008 R2, and only for use with the Hyper-V role. A CSV may be created only from shared storage, such as a LUN provisioned from the NetApp storage array, to which all Hyper-V server nodes in the failover cluster are connected.

A CSV functions as a distributed-access file system that is optimized for Hyper-V and uses a standard New Technology File System (NTFS). This means that any shared storage supported as a standard cluster volume can be configured as a CSV. CSVs enable multiple Hyper-V server nodes to concurrently access the same disk or LUN.

The introduction of CSVs in Windows Server 2008 R2 Hyper-V has many advantages, including the following:

- **Simplified storage management.** More VMs share fewer LUNs.
- **Independent movement of VMs.** While multiple clustered VMs share the same LUN, they are still able to migrate or fail over between Hyper-V cluster nodes independently of one another.
- **Shared namespace.** CSVs do not need to be assigned a drive letter, reducing restrictions on such and eliminating the need to manage GUIDs or mount points.
- **Enhanced availability.** CSVs are able to detect and address additional failure events that would otherwise cause the VMs to be unavailable.
- **Storage efficiency.** Pooling VMs on the same LUN simplifies capacity planning and reduces the amount of space reserved for future growth, as it is no longer set aside on a per-VM basis.

CSVs are implemented as a file system minifilter, within `CSVfilter.sys`, as shown in Figure 6, and the minifilter is responsible for intercepting NTFS metadata requests and all of the I/O requests.

Figure 6) CSV is implemented as a file system minifilter.

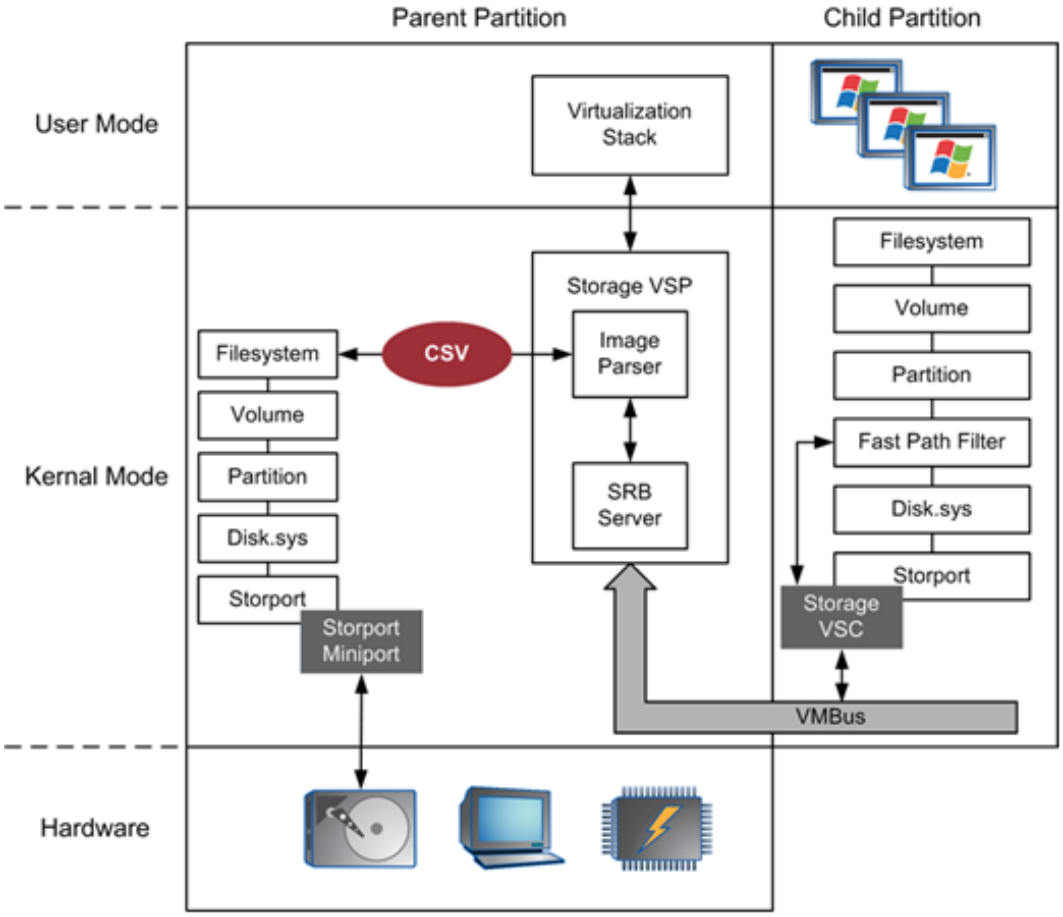

The Hyper-V cluster node where the LUN was first enabled as a CSV becomes the CSV owner node (also the coordinator node) for the CSV. Although all Hyper-V cluster nodes can access the CSV concurrently, the CSV nonowner nodes must forward all metadata operations (such as changes to file size or properties) to the CSV owner node, but all nodes can perform data I/O-level operations on the CSV; this is shown in Figure 7.

Figure 7) CSV metadata and data I/O operations between Hyper-V cluster nodes.

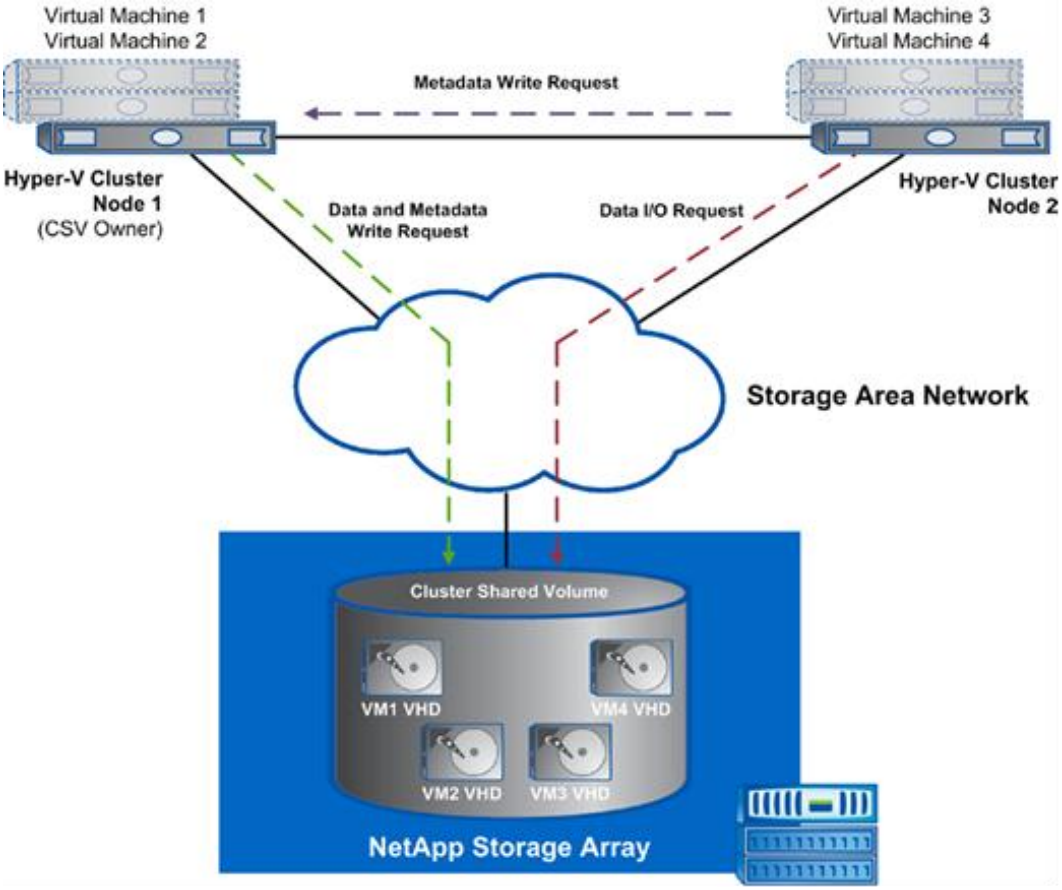

Cluster Shared Volumes Single Namespace

A feature of enabling a standard cluster volume to become a CSV is the availability of a single namespace between all Hyper-V server nodes in the failover cluster, where any files (such as VHDs) have the same file path and name from any Hyper-V server node in the failover cluster. All CSVs are stored as directories and subdirectories beneath the `%systemdrive%\ClusterStorage` root folder, as illustrated in Figure 8.

Figure 8) Cluster shared volume single namespace.

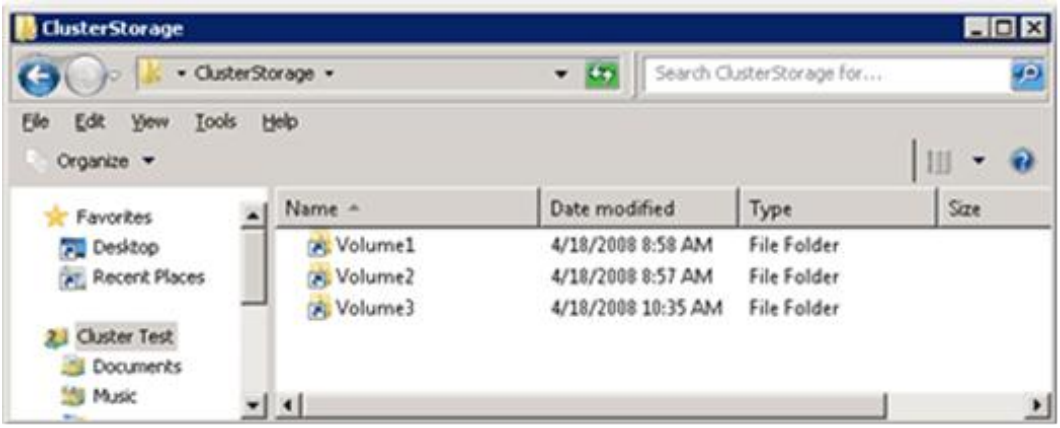

As Figure 8 shows, the CSVs (Volume1, Volume2, and Volume3) are located in the ClusterStorage folder. If the system drive were C:\, then the fully qualified paths to the earlier mentioned volumes would be:

- `C:\ClusterStorage\Volume1`
- `C:\ClusterStorage\Volume2`
- `C:\ClusterStorage\Volume3`

This means that each CSV does not require a drive letter reserved for it, nor do administrators need to manage volume mount points or GUIDs. This reduces restrictions on drive letter assignments and simplifies management of the CSVs.

Live Migration

As with a standard cluster volume in Windows Server 2008 R2 Hyper-V, live migration is supported with CSVs as well. However, there is an advantage to deploying CSVs in the Hyper-V environment when it comes to live migration, especially overuse of standard cluster volumes. When you are live migrating VMs that reside on a CSV, using CSVs reduces the brown-out period at the end of the live migration process. This is because a standard cluster volume would have to be unmounted and remounted as part of the live migration process, introducing additional delay that could affect the availability of the VMs, but this is not the case with CSVs because they are seen from all Hyper-V cluster nodes.

Virtual Machine Architecture

Because a CSV allows concurrent access to VMs and VHDs between all Hyper-V cluster nodes, Microsoft no longer restricts or recommends deployment of “one VM per LUN.” Therefore, an administrator may place any number of VMs and VHDs on the CSV, as long as the CSV’s underlying storage is capable of handling the performance requirements of the VMs and VHDs located within it.

Virtual machines that are deployed with pass-through disks are supported with live migration, but pass-through disks cannot be enabled as CSVs. However, you might choose to keep the VM configuration files on a CSV or use a VHD for the system volume and store that on a CSV while using pass-through disks for all other VM storage.

Dynamic I/O Redirection

As the result of improvements to Windows failover clustering with Windows Server 2008 R2, administrators will benefit from improved fault tolerance across the Hyper-V server nodes in the cluster. With CSV comes a feature called dynamic I/O redirection, which allows both storage and network I/O to be redirected within the failover cluster based on connection availability. It is depicted in Figure 9.

After I/O for the CSV is directed, it reports the event to the Windows event logs, and it shows within Failover Cluster Manager as well, as Figure 10 shows.

As might be expected, when storage I/O is redirected over the network, especially where bandwidth of the IP network might be smaller than the original failed storage path, performance of the Hyper-V cluster node, its storage subsystem, and VMs could be affected. An administrator can do a few things to minimize the opportunity for events that cause I/O redirection to occur in the first place, as well as minimize any interruption to operation of the Hyper-V cluster nodes or VMs as the result of I/O redirection occurring.

Best Practices

- **Networks**: In addition to the NICs installed in the Hyper-V server for management, VMs, IP storage, and more, NetApp strongly recommends that you dedicate a physical network adapter to CSV traffic only. The physical network adapter should be a minimum of 1GB. If you are running large servers (16 LCPUs+, 64GB+), planning to use CSVs extensively (many, large numbers of), or planning to balance VMs dynamically across the cluster using SCVMM or just to use live migration itself extensively, you should consider using 10GB Ethernet for CSV traffic instead.
- **Storage**: NetApp strongly recommends that you configure MPIO on all Hyper-V cluster nodes to minimize the opportunity for CSV I/O redirection to occur. More information on configuring MPIO with NetApp storage can be found in "Multipathing/MPIO" in section 5.1. Also, for the latest supported software versions and related requirements, see the NetApp Interoperability Matrix.

Microsoft makes additional recommendations in Hyper-V: Using Hyper-V and Failover Clustering and Hyper-V: Using Live Migration with Cluster Shared Volumes in Windows Server 2008 R2, which NetApp supports, and therefore also recommends that you adopt all best practices described in these documents in the interest of supportability and maximum performance.

Figure 9) Cluster shared volumes dynamic I/O redirection.

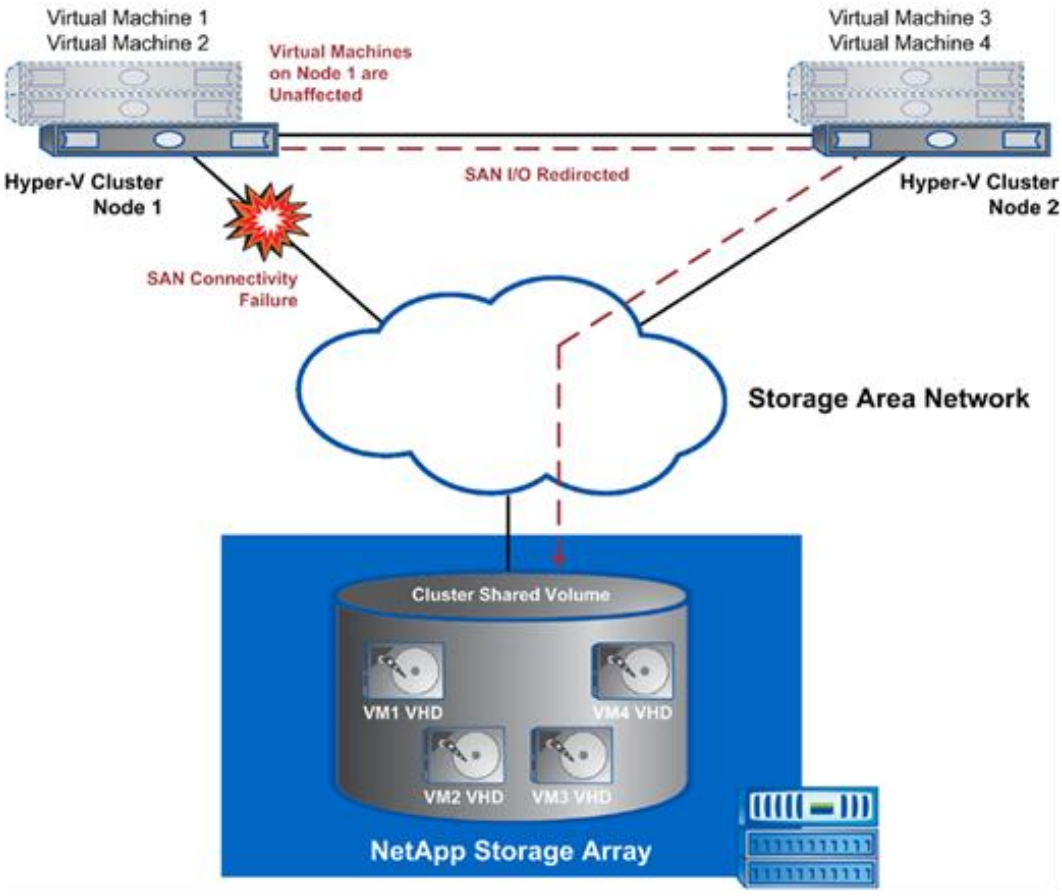

Figure 10) CSV I/O redirection shows as "Redirected Access" in the Failover Cluster Manager MMC.

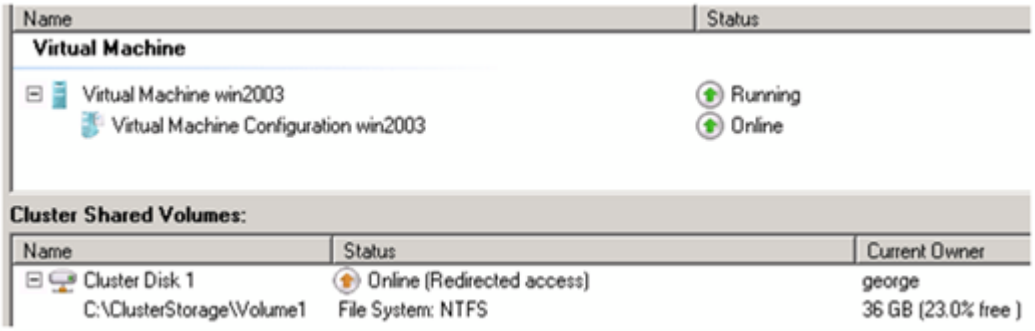

CSV Failure Events

The first type of failure to which CSV responds is the loss of connectivity between a Hyper-V server node and the shared storage. Because Hyper-V supports both iSCSI and Fibre Channel connectivity to shared storage, CSV respond to the loss of connectivity on an IP network and Fibre Channel network. As shown in Figure 11, if the SAN connection from Hyper-V server node 1 fails, the I/O operations are redirected over the IP network to Hyper-V server node 2. Hyper-V server node 2 then handles the I/O operations to the SAN for Hyper-V server node 1.

Figure 11) CSV I/O redirection in the event of a storage path failure.

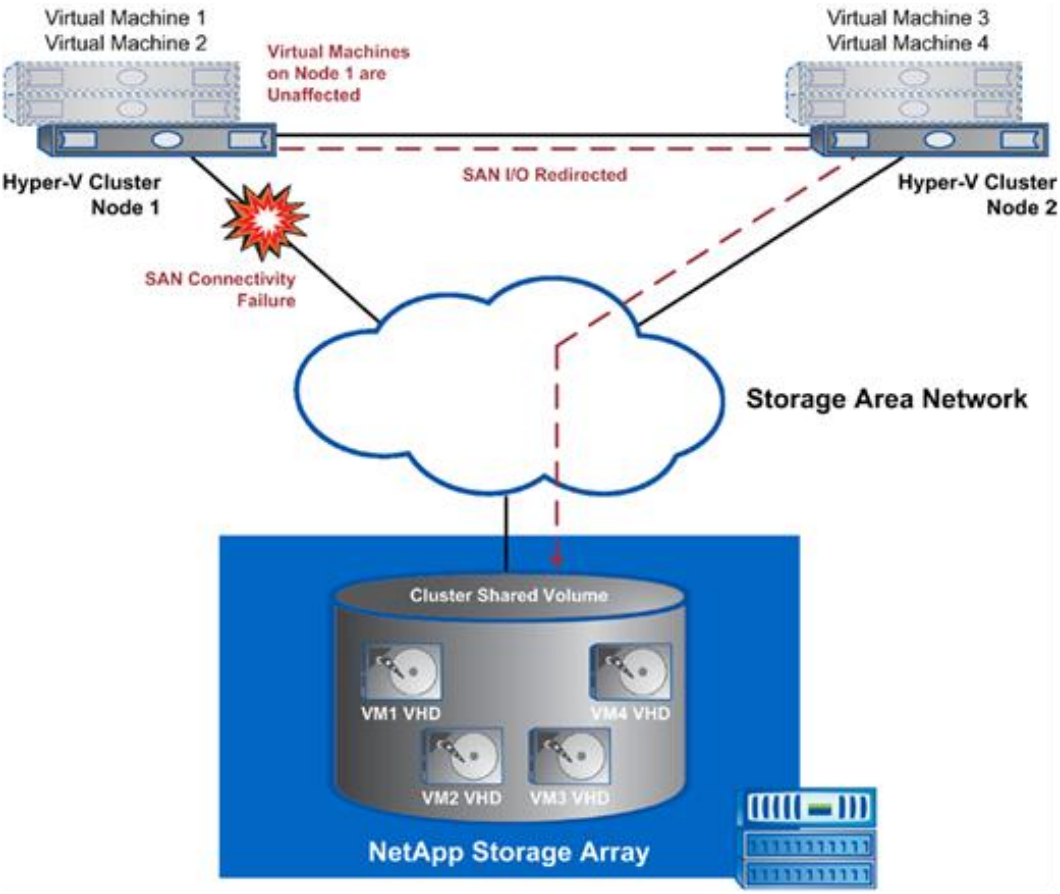

Best Practice

After the administrator has recognized that CSV I/O redirection is occurring for a specific CSV, the administrator should immediately live migrate all VMs residing on that CSV to another node.

The second type of failure to which CSV responds is the failure of a network path. As Figure 12 shows, if the cluster detects that a network path has been lost, then it fails the network traffic from that path over to the next available path on that Hyper-V cluster node and finally to the path used for CSV communication.

The third type of failure to which CSV responds is the failure of an entire Hyper-V cluster node. Even if the Hyper-V cluster node is the owner of some CSVs, it transfers ownership of that CSV to another Hyper-V cluster node (Figure 13).

Figure 12) CSV I/O redirection in the event of a network path failure.

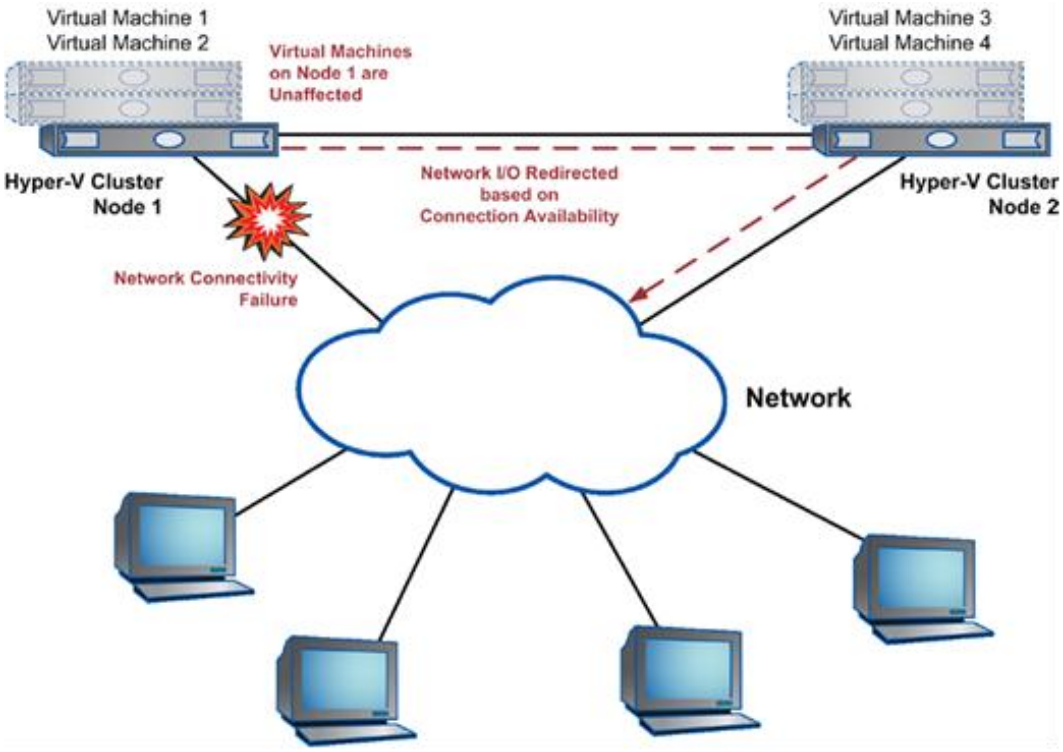

Figure 13) Volume ownership transferred in the event of a Hyper-V cluster node failure.

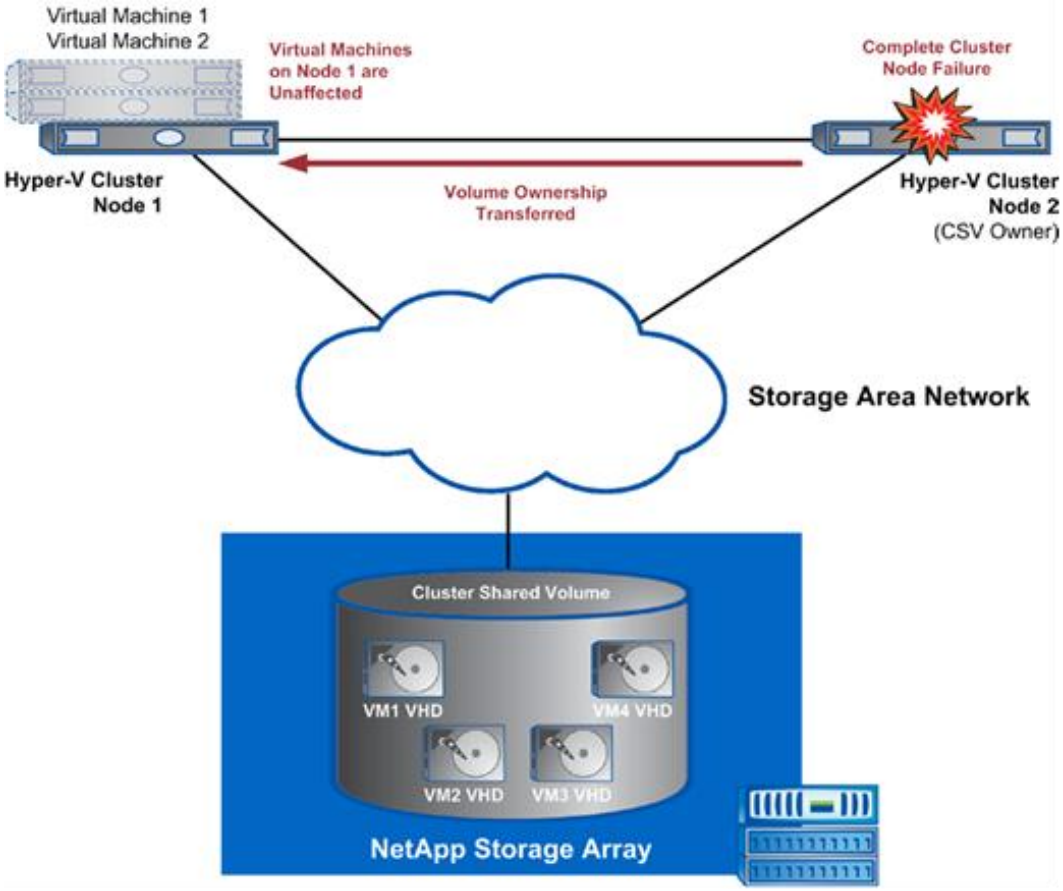

Direct I/O

As discussed earlier in the beginning of Cluster Shared Volume (CSV) in section 5.3, although all Hyper-V cluster nodes can access the CSV concurrently, the CSV nonowner nodes must forward all metadata operations (such as changes to file size or properties) to the CSV owner node, but all nodes can perform data I/O–level operations on the CSV (Figure 7). However, for VM I/O operations on CSV nonowner nodes, data access on the CSV is conducted through a new technique called direct I/O.

Direct I/O refers to the process used with CSV-enabled volumes that allow I/O operations from a VM to be performed on the storage device with enhanced performance. This is accomplished by splitting the I/O transaction and sending the NTFS metadata operations to the CSV owner node over the network while sending the I/O directly to the disk class driver for the storage device. By sending data I/O directly from the file system minidriver (`CSVFilter.sys`) to the disk class driver (`disk.sys`), CSV bypasses all local file system and volume/partition processing on nonowner nodes, as shown in Figure 14.

Figure 14) Direct I/O bypasses file system/volume/partition processing on CSV nonowner nodes only.

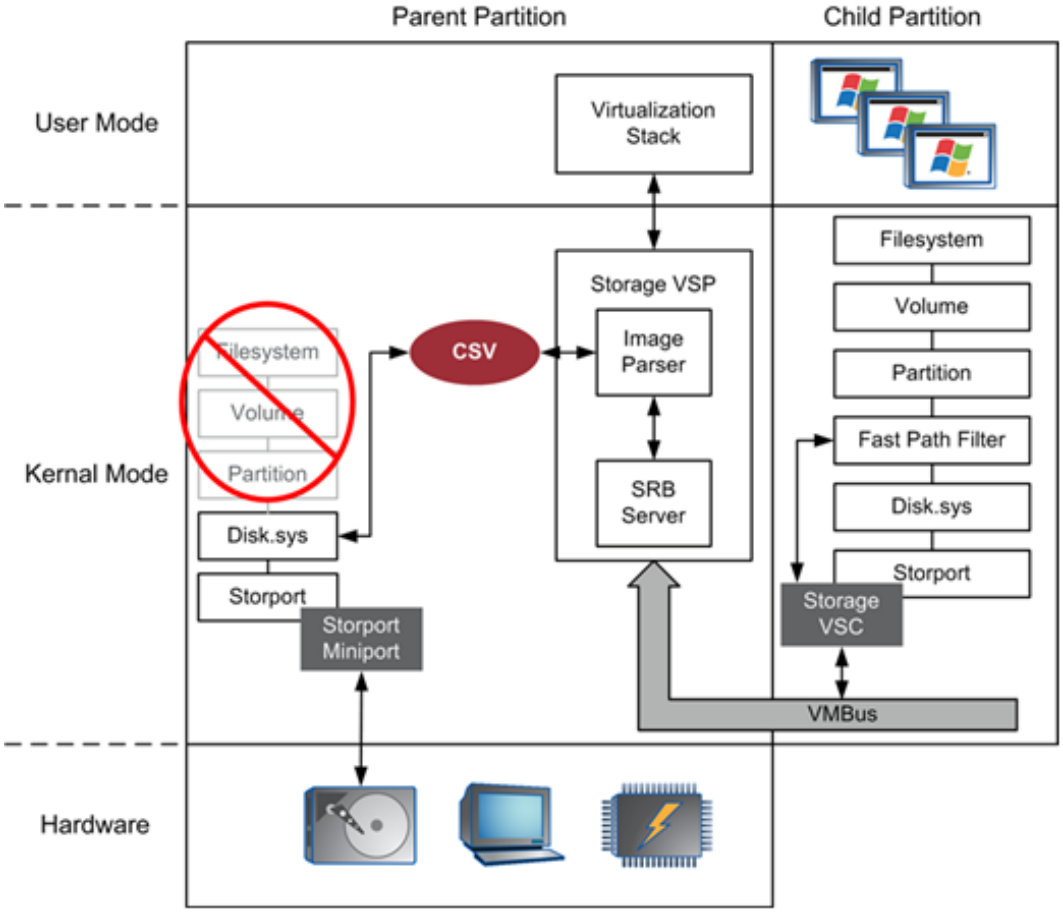

Backup and Recovery

Backup and recovery of CSVs in a Hyper-V environment can be tricky. Microsoft has released information stating that all vendors must make changes in their products to support backup and recovery of VMs residing on CSVs and the CSVs themselves.

Why do backup vendors have to make changes to support CSVs? The architecture of CSVs is complex, so we will avoid discussing it in length. However, the CSV architecture introduces minute changes to how data is handled in a failover cluster environment. As discussed in section 5.3, under “Cluster Shared Volume (CSV)”, when it comes to backups, the owner node becomes the key to obtaining consistent backups of the VMs residing on a CSV. While the CSV is normally accessible for data I/O operations by all nodes in the cluster, regardless of whether an owner or a nonowner node, the owner node is the single node in the cluster that has access to the Volume Shadow Service (VSS) components. Because the VSS infrastructure is present only on the owner node, this is the only node where backups of VMs can be obtained in a consistent manner. Therefore, the CSV is transferred between Hyper-V cluster nodes as needed to back up VMs as necessary. CSV ownership might change multiple times during a normal backup window to make this possible. The backup application must be aware of these changes and know whether it’s possible to obtain a VSS-integrated backup of certain VMs.

NetApp SnapManager for Hyper-V assists administrators with backup and recovery of virtual machines in a Hyper-V environment and is currently available for download through the NetApp Support site. SnapManager for Hyper-V (SMHV) supports backup and recovery of VMs residing on CSVs using NetApp Snapshot technology. SMHV can provide application- and data-consistent Snapshot copies of

virtual machines. This includes not just the guest OS, but also any application installed that is compatible with VSS. SnapManager for Hyper-V is discussed in Chapter 12 and following in this report.

Best Practices

Because all backups are coordinated and occur on the CSV owner node, NetApp recommends that administrators put some additional thought into how VMs and CSVs are provisioned among all the Hyper-V servers.

In some environments, administrators might consider provisioning a greater number of smaller CSVs or a smaller number of larger CSVs to their Hyper-V cluster. Although this is the administrator's choice, customers who consider the latter choice should do so only when performance of the underlying disk from the NetApp storage system can support a larger number of VMs.

Regardless of choice with sizing and provisioning CSVs within the Hyper-V cluster, NetApp recommends balancing the VMs across the deployed CSVs as well as possible. By organizing the VMs deployed to the CSVs, you balance the backup operations that occur at similar times across all CSVs.

Irrespective of the considerations mentioned earlier, NetApp asks administrators to consider designing for performance first when sizing and provisioning CSVs for Hyper-V servers.

Virtual Machine Provisioning

Provisioning virtual machines in a Hyper-V environment in which CSVs are deployed can also require some forethought. For the same reason that there are changes to backup and recovery with CSVs, there is also reason to consider changing the way you deploy VMs to CSVs. Because of how I/O and metadata operations are conducted on Hyper-V cluster nodes, depending upon whether they are owner or nonowner nodes, it should be expected that performance between an owner and a nonowner node can differ in certain environments.

Microsoft System Center Virtual Machine Manager (SCVMM) 2008 R2 recognizes the differences between an owner and nonowner node when deploying VMs to the Hyper-V cluster nodes. When deploying a VM, it creates the VHD on the CSV by conducting all operations through the Hyper-V CSV owner node. Essentially, the VHD is deployed to the Hyper-V CSV owner node first, while the VM is created on the host selected during the VM creation wizard within SCVMM. Since CSVs have a single namespace, the path to the VHD is the same regardless of the Hyper-V server to which it was originally deployed.

The fact that Microsoft has designed SCVMM to behave this way when deploying VMs and VHDs to CSVs gives an indication about how you should consider deploying VMs and VHDs in your Hyper-V environment, especially if you are not using SCVMM.

Best Practice

When using Hyper-V Manager to create a new VM, first create the VHD on the appropriate CSV volume and do so from the Hyper-V CSV owner node. When using Failover Cluster Manager to create a new VM, first create the VHD by selecting the Hyper-V CSV owner node. After the VHD is created, then create the VM on any Hyper-V cluster node and select the previously created VHD as part of the creation process.

If you are using SCVMM to deploy your VMs, then no additional action is necessary.

5.4 Virtual Machine Storage Provisioning

Microsoft Hyper-V virtual machines have the option to store their data on these devices:

- Virtual hard disks (VHDs)

- Pass-through disks
- Directly connected LUNs using the Microsoft iSCSI Software Initiator

With all three of these storage options, the disks are formatted with the file system of the child OS.

Virtual Hard Disks (VHDs)

One option for configuring virtual disks for VMs with Microsoft Hyper-V is through the use of virtual hard disks or VHDs. VHDs allow you to assign storage to a VM while the actual storage is kept in a VHD file located on a physical disk, which is formatted with NTFS and attached to the Hyper-V parent partition. Often these physical disks are LUNs configured from shared storage, through Fibre Channel or iSCSI only. VHDs offer the benefits of increased manageability and portability associated with having the VM storage encapsulated in a single file.

There are three different types of VHDs:

- Dynamically expanding VHD
- Fixed-size VHD
- Differencing VHD

Fixed-Size VHD

This type of VHD allocates the entire configured amount of storage at the time of creation; thus, it does not expand while the VM is operational under any circumstances. Therefore, the fixed-size VHD type offers the lowest performance overhead when compared to the other Hyper-V VHD types. The upfront data storage requirements of fixed-size VHDs can be mitigated by many NetApp technologies that help regain storage savings, such as thin provisioning and deduplication (for more information, see Chapter 6, “Increasing Storage Efficiency and Flexibility”).

Best Practice
NetApp strongly recommends using fixed-size VHDs to achieve best performance for your Microsoft virtual environment. Minimize use of dynamically expanding and differencing VHDs in a production environment because of their higher performance overhead unless you have a solid reason to consider a configuration using these types of VHDs.

Dynamically Expanding VHD

This type of VHD doesn’t allocate the entire configured amount of storage at the time of creation; thus, it expands while the VM is operational and data is being written to the virtual disk, but it never expands beyond its configured size. Primarily, dynamically expanding VHDs differ in the area of performance, because they have the highest performance overhead when compared to the other Hyper-V VHD types. The biggest impact to performance comes from the expansion of the VHD while the VM is operational and data is being written to the virtual disk. In addition, how Hyper-V expands the VHD during operation has been observed to introduce file system alignment issues and can also cause fragmentation to the child file system. Therefore, the child OS must be adequately monitored and managed in order to prevent performance degradation of the child. However, there might be reasons to give further consideration to use dynamically expanding VHDs, such as test and development environments or scenarios in which you expect the child to expand its data storage needs ad hoc.

Differencing VHD

This type of VHD is not created at the time the VM is initially configured, but rather when a Hyper-V snapshot is taken of an existing VM. A differencing VHD points to a parent VHD file, which can be any type of VHD and functions similarly to a dynamically expanding VHD. Because of their similarities to dynamically expanding VHDs, differencing VHDs suffer from the same performance concerns, and

therefore you should review the same considerations. Note that Hyper-V snapshots have no relation to NetApp Snapshot technology.

Pass-Through Disks

Pass-through disks are disks that are attached directly to the Hyper-V parent but assigned directly to a VM and formatted with the child OS file system. A pass-through disk is best suited for large datasets (typically beyond 300GB to 500GB) and extreme I/O requirements. One of the limitations associated with the use of pass-through disks with VMs is that Hyper-V snapshots are not supported. For this reason, NetApp recommends limiting use of pass-through disks in your Hyper-V environment, except where considered necessary.

Best Practice

While NetApp strongly recommends fixed-size VHDs for the majority of VM configurations, they have some limitations. Therefore, NetApp recommends use of pass-through disks for VMs with large datasets, typically beyond 300GB to 500GB, with extreme I/O requirements, and for those virtual disks requiring more than 2TB of storage.

Storage Presented Directly to Child OS

In addition to utilizing VHDs and pass-through disks to configure additional storage for VMs, you might also configure storage directly within the child OS by using an iSCSI software initiator supported by the child OS.

This allows you to provision storage directly to the VM for installation of applications and storage of application data. Being able to use NetApp SnapDrive for Windows to provision the storage to the VM allows you to manage applications with the NetApp SnapManager family of products. NetApp supports all of the SnapManager products for Microsoft Exchange, SharePoint® Server, SQL Server®, as well as SAP® and Oracle® running within a Hyper-V VM.

Currently, Microsoft does not support booting a child OS over iSCSI that is mapped using the network into the VM. If booting from an iSCSI LUN is desired for a child OS, then the iSCSI LUN should be configured to the Hyper-V parent as a pass-through disk and mapped into the VM.

Best Practice

For NetApp products such as SnapManager to work within guest OSs, and for support of advanced data protection configurations with SnapMirror, having LUNs provisioned directly to the guest OS is required.

Virtual Disk Devices: IDE Versus SCSI

As Table 12 shows, there are two options for presenting or exposing storage to the Hyper-V VM when using VHDs or pass-through disks: IDE or SCSI device. While the first virtual disk presented to a VM must use an IDE device, an administrator might choose between an IDE or SCSI device when presenting all additional storage to the VM.

Table 12) Hyper-V storage comparison table.

	DAS or SAN on Host, VHD, or Pass-Through Disk on Host, Exposed to VM as IDE	DAS or SAN on Host, VHD or Pass-Through Disk on Host, Exposed to VM as SCSI	Not Exposed to Host, Exposed to Child as Directly Connected iSCSI LUN
Child boot from disk	Yes	No	No

	DAS or SAN on Host, VHD, or Pass-Through Disk on Host, Exposed to VM as IDE	DAS or SAN on Host, VHD or Pass-Through Disk on Host, Exposed to VM as SCSI	Not Exposed to Host, Exposed to Child as Directly Connected iSCSI LUN
Additional SW on child	Integration components (optional)	Integration components	iSCSI SW Initiator
Child sees disk as	Virtual HS ATA device	Microsoft virtual disk SCSI disk device	NetApp LUN SCSI disk device
Child maximum disks	2 x 2 = 4 disks	4 x 64 = 256 disks	Not limited by Hyper-V
Child hot-add disk	No	Yes with R2 only	Yes

Table 12 also shows significant differences between choosing an IDE or SCSI device when presenting storage to the VM. Often the most important considerations an administrator has for choosing between an IDE or SCSI device are the option to install Integration Services, the ability to later hot add/remove storage to the VM, and the maximum number of disks that will be presented to the VM.

In order to take advantage of the ability to hot add/remove storage to the VM with Hyper-V R2 (not available prior to this release), as well as to support the maximum number of virtual disks per VM, an administrator must make that storage available to the VM using a SCSI device. However, if an administrator wants to use a SCSI device to present additional storage to the VM, the child OS must have Integration Services installed.

Best Practice

For maximum performance, use SCSI devices to present virtual disks to the VMs when possible.

Integration Services

As discussed in the preceding paragraphs, installation of Integration Services within a guest OS is an important consideration, not just in choosing between an IDE or a SCSI device when presenting storage to a VM, but also when considering the performance requirements of the VM.

The integration components install “enlightened” drivers, among other things, to optimize the overall performance of a VM. Specifically, these enlightened drivers provide support for the synthetic I/O devices, which significantly reduces CPU overhead for I/O when compared to using emulated I/O devices and allows the synthetic I/O device to take advantage of the unique Hyper-V architecture not available to emulated I/O devices. If the child OS is not supported by Hyper-V R2, then the VM is forced to use emulated devices, because there is not an option to install Integration Services within a VM that has an unsupported guest OS. For more information on supported guest OSs with Hyper-V R2, see Virtualization with Hyper-V: Supported Guest Operating Systems.

Best Practice

For maximum performance, make sure that Integration Services are installed within all child OSs where supported. In accordance with Microsoft’s own best practices, the version of Integration Services installed within child OSs should always correspond directly to the OS version of the Hyper-V parent partition. Where this is not so, such as during upgrades of the underlying Hyper-V host or migration of VMs between Hyper-V parent partitions with different versions, it is supported only in nonproduction modes.

Virtual Machine Disk Performance

Microsoft has published much information on the performance of the different virtual machine disk options, both on Hyper-V version 1 and Hyper-V R2, especially since Microsoft claims that it has radically improved the performance of its virtual disk options with Hyper-V R2, which has brought the expected performance of VHDs and pass-through disks much closer to native. For more information on the performance characteristics of VHDs and pass-through disks, as well as the performance of the underlying Hyper-V storage architecture, see some of the resources listed here.

- Microsoft Windows Server 2008 Hyper-V R2
  - Performance Tuning Guidelines for Windows Server 2008 R2 (specifically the section titled "Storage I/O Performance")
  - All Topics Performance: What's New in Windows Server 2008 R2 Hyper-V Performance and Scale? (specifically the section titled "Storage improvements from Windows Server 2008 SP1 Hyper-V to Windows Server 2008 R2 Hyper-V")
  - Microsoft Virtual Hard Disk (VHD) FAQ and Virtual Hard Disk (VHD) Image Format Specification
- Microsoft Windows Server Hyper-V
  - Performance Tuning Guidelines for Windows Server 2008 (specifically the section titled "Storage I/O Performance")
  - Windows Server Performance Team Blog: Hyper-V and VHD Performance: Dynamic vs Fixed
  - Microsoft Virtual Hard Disk (VHD) FAQ and Virtual Hard Disk (VHD) Image Format Specification

Virtual Machine Sizing

Before you begin to provision LUNs, you must consider the storage requirements for each VM in order to plan and correctly provision LUNs that will support the VMs stored on them (VHDs) or VMs using them (pass-through disks). This section is designed to help administrators understand the space requirements for VMs to be deployed in the Hyper-V environment. Consider that the recommendations in this section best apply to production environments using shared storage and that your mileage might differ slightly in other environments.

When determining the appropriate storage to provision for use with VMs, there are a few considerations for administrators to keep in mind:

- VMs can generally be classified into three different categories, those requiring high, medium, and low disk utilization. For VMs that are classified as high, an administrator might have to consider isolating the VM and its data on its own LUNs.
  - Where VHDs are being used, consider provisioning multiple VHDs to the VM, separating system data, application data, application logs, system page file, and so forth on separate VHDs. These VHDs can be isolated on their own LUNs or for some data, such as VHDs containing the system page file, can be combined with other VHDs belonging to other VMs.
  - Either exclusively or in combination with use of VHDs, a pass-through disk might be appropriate for storing some data, such as application data or logs, effectively isolating that data on its own LUN. However, caution should be exercised with use of pass-through LUNs as they make migration of VMs and backup of VMs more complicated.
  - As much as possible, with the exception of the system data, use the SCSI disk device to connect each VHD or pass-through disk. Because Hyper-V supports a maximum of four SCSI disk controllers per VM, some SCSI controllers might require more than one VHD or pass-through disk attached to each one.
- In some configurations, especially where CSVs are used, multiple VMs and VHDs are combined on the same LUN. Therefore, it might not be as simple as calculating the storage requirements for a single VM when multiple VMs are located on the same LUN. This is especially true when separating VM data types across multiple shared LUNs, such as a single CSV to store all VM configuration files or VHDs containing VM system data.

- For VMs using dynamically expanding VHDs, administrators must correctly size the LUN to match the maximum size of the VHD or monitor the LUN extremely carefully.

The information in Table 13 is meant to help administrators understand the space requirements for each VM, rather than to size actual LUNs. This table should be used for each VM in your virtual environment, and then the information for all VMs should be aggregated to size your LUNs appropriately based on the number of VMs and the types of data you plan to store on that LUN.

Table 13) Virtual machine storage sizing worksheet.

VM Component	Component Size	Component Description
VM configuration	200MB	Includes VM configuration and VM state (XML and VSV) files
VM memory		Includes the VM memory (BIN) file
Guest OS system data		Depending on guest OS type, the space required will differ but will always reside in the IDE device of the VM.
Guest OS page file		Microsoft recommends 1.5 times the amount of memory allocated to the VM. This can reside with the system data or separately, but it should be calculated separately.
Guest OS application data		Additional storage is required for application installations, data, log files, and more.
VM Hyper-V snapshots	10–15% times guest OS system data page file, and application data	Although NetApp recommends using NetApp Snapshot copies whenever possible because they are more storage efficient and simplify storage planning, if Hyper-V snapshots will be used, space must be planned for them.

Always size a LUN larger than the actual storage required to allow room for data growth because VMs might require more memory, expansion to storage required for VHDs, and room for Hyper-V snapshots. How much larger should be decided on a case-by-case basis based on the actual VMs stored and VM data types stored. However, LUNs can be resized or shrunk on the fly, with no downtime, when using NetApp SnapDrive and NetApp System Manager.

Virtual Machine Disk Alignment

For optimal performance, the file system of the VHD, the underlying physical disk on the Hyper-V server, and the storage array should be in proper alignment.

Best Practice
Failure to make sure that all VM disks are aligned can result in mild to severe performance degradation, depending upon the configuration and the overall environment.

File System Alignment Concepts

Historically, hard drives (LUNs) presented the OS with a logical geometry that would be used to partition and format the disk in an efficient manner. Logical geometry today is virtual and fabricated by the host BIOS and operating system. Operating partitioning programs such as fdisk use the fake disk geometry to determine where to begin a partition. Unfortunately, some partitioning programs create disk partitions that do not align to underlying block boundaries of the disk. Notable examples of this are the GNU fdisk found on many Linux® distributions and Microsoft Diskpart found on Windows 2000 and Windows 2003.

NetApp uses 4KB blocks (4 x 1,024 = 4,096 bytes) as its basic storage building block. Write operations can consume no less than a single 4KB block and can consume many 4KB blocks, depending on the size

of the write operation. Ideally, the guest/child OS should align its file systems so that writes are aligned to the storage device's logical blocks. The problem of unaligned LUN I/O occurs when the partitioning scheme used by the host OS doesn't match the block boundaries inside the LUN, as shown in Figure 15. If the guest file system is not aligned, it might become necessary to read or write twice as many blocks of storage than the guest actually requested since any guest file system block actually occupies at least two partial storage blocks. As a simple example, assuming only one layer of file system and that the guest allocation unit is equal to the storage logical block size (4K or 4,096 bytes), each guest block (technically an "allocation unit") would occupy 512 bytes of one block and 3,584 bytes (4,096–512) of the next.

This results in inefficient I/O because the storage controller must perform additional work such as reading additional data to satisfy a read or write I/O from the host.

By default, many guest operating systems, including Windows 2000 and 2003 and various Linux distributions, start the first primary partition at sector (logical block) 63. The reasons for this are historically tied to disk geometry. This behavior leads to misaligned file systems since the partition does not begin at a sector that is a multiple of eight.

Figure 15) Misaligned file system.

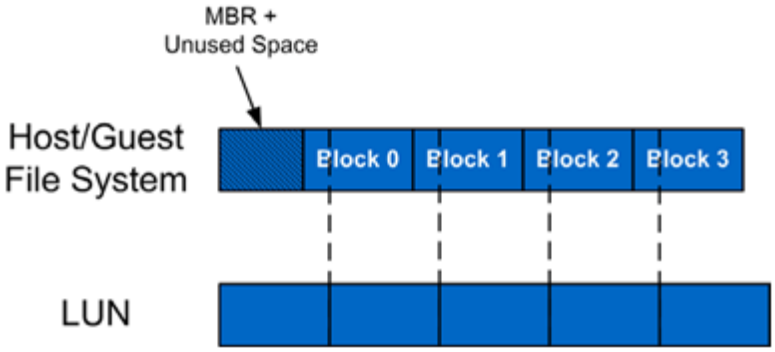

This issue is more complex when the file system on the virtualization host contains the files (vmdk, vhd) that represent the VM virtual disks, as shown in Figure 16 and Figure 17. In this case, the partition scheme used by the guest OS for the virtual disks must match the partition scheme used by the LUNs on the hypervisor host and the NetApp storage array blocks.

Figure 16) Guest OS and NTFS file system are not aligned with the NetApp storage array blocks.

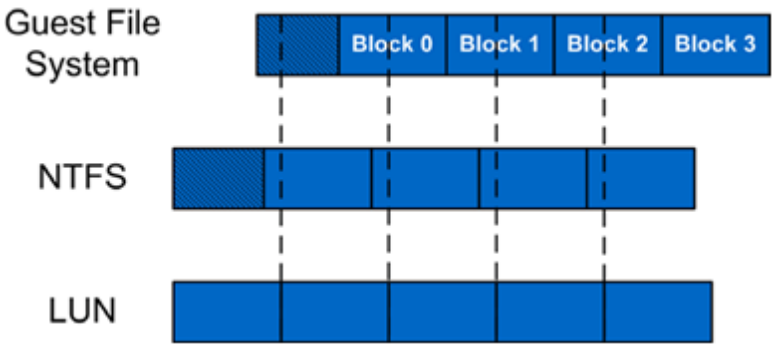

Figure 17) Guest OS and NTFS file system are aligned with the NetApp storage array blocks.

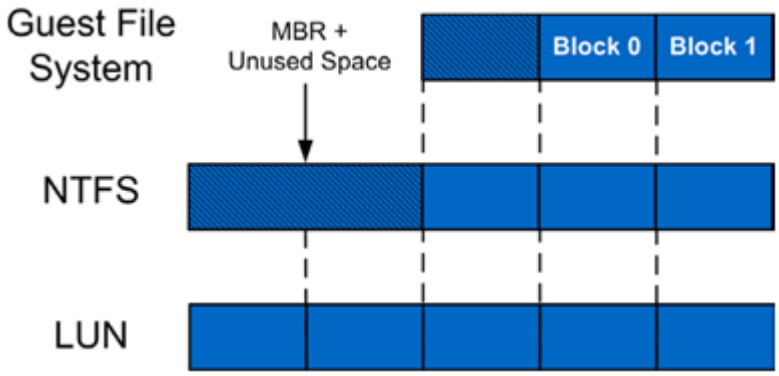

For VHDs hosted on NTFS formatted LUNs attached as physical disks on the Hyper-V parent partition, two layers of alignment are involved. The NTFS file system on the physical disk and the file system on the child VM hosted on the physical disk should align with the NetApp storage blocks.

Pass-through disks and LUNs directly mapped by the child VM do not require special attention if the LUN type of the LUN matches the guest operating system type. For more information on LUN type, see the "LUN Multiprotocol Type" section in the Data ONTAP Block Access Management Guide or the Commands Manual Page Reference Document, Volume 1 and Volume 2, which can be downloaded from the NetApp Support site.

Table 14) Layers of storage for file system alignment with Microsoft Hyper-V.

Layers of Storage	Hyper-V Shared Storage Options		
	NTFS-Formatted LUNs	Pass-Through Disks	LUNs Directly Mapped to the Child OS
Child OS	✓	✓	✓
Hyper-V parent partition	✓	N/A	N/A
NetApp storage array	✓	✓	✓

The different layers of storage involved for each of the shared storage options are shown earlier in Table 4 through Table 6. The check mark (✓) indicates that alignment should be enforced at the guest OS, Hyper-V parent partition, and/or NetApp storage array level.

Enforcing VM disk alignment starts with selecting the correct LUN protocol type on the NetApp storage array. This way, once the file system on the physical disk, the pass-through disk, or the directly connected disk is formatted, alignment to the NetApp storage array can be assured. For details, see "LUNs" in section 5.2 and follow the instructions in "Selecting the Correct LUN Protocol Type."

For existing child VMs that are misaligned, NetApp recommends correcting the offset of only child VMs that are experiencing an I/O performance issue. This performance penalty should be more noticeable on systems that are completing a large number of small read and write operations. The reason for this recommendation is that in order to correct the partition offset, a new physical disk must be created and formatted, and the data must be migrated from the original disk to the new one.

Best Practice

NetApp strongly recommends correcting the offset for all VM templates and existing VMs that are misaligned and experiencing an I/O performance issue. Misaligned VMs with low I/O requirements might not benefit from the effort to correct the misalignment.

File System Alignment Prevention

Virtual Hard Disks (VHDs)

For VHDs hosted on NTFS-formatted LUNs attached as physical disks on the Hyper-V parent partition, two layers of alignment are involved. The NTFS file system on the physical disk and the file system on the child VM hosted on the physical disk should align with the NetApp storage blocks. For details, see Figure 17 and Table 14.

Proper file system alignment can be assured only with fixed-size VHDs. With dynamically expanding and differencing VHDs, there is no guarantee of proper file system alignment, and there is a possibility of a performance penalty as well.

Best Practice

NetApp recommends using fixed-size VHDs within your Microsoft Hyper-V virtual environment wherever possible, especially in production environments. Proper file system alignment can be assured only on fixed-size VHDs.

Avoid use of dynamically expanding and differencing VHDs whenever possible because file system alignment can never be guaranteed with these VHD types.

Aligning the file system of the VHD to the file system of the underlying physical disk provides the best possible performance as the system scales out. All VHD types can be formatted with the correct offset at the time of creation by booting the child VM before installing an OS and manually setting the partition offset. For Windows child VMs, one might consider using the Windows Preinstall Environment boot CD or alternative tools such as Bart's PE CD.

When the VHDs are aligned for use with NetApp storage systems, the starting partition offset must be divisible by 4,096. The recommended starting offset value for Windows OSs is 32,768. For a Windows child OS, verifying this value is easy by using the msinfo32 utility. The default starting offset value typically observed is 32,256. For details, see "File System Alignment Correction" later in this section, specifically the subsection titled "Detection."

Pass-Through Virtual Disks

Figure 18) Child file system aligned with the storage array blocks.

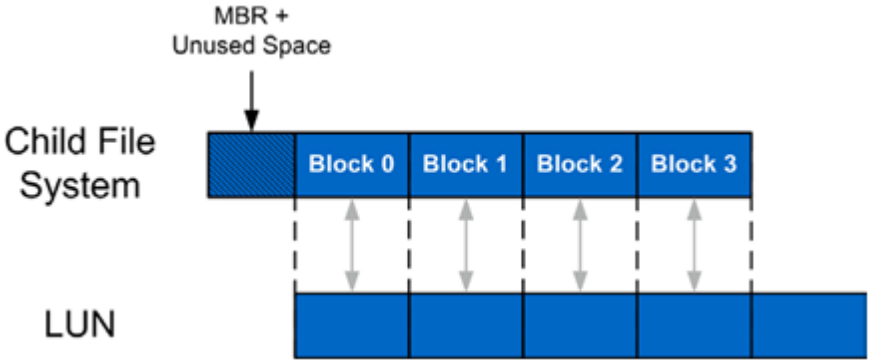

Pass-through disks do not require special attention if the LUN protocol type of the LUN matches the guest operating system type. For details, see the LUNs section in this report and follow the instructions in "Selecting the Correct LUN Protocol Type." In addition, see the LUN Multiprotocol Type section in the Data ONTAP Block Access Management Guide for iSCSI or FC or Data ONTAP Commands Manual Page Reference Document, Volume 1 or Volume 2, on the NetApp Support site.

Disks Attached Directly to the Child OS

For LUNs attached directly to the child OS, only one layer of alignment is involved, similar to pass-through disks. Only the child OS must align with the NetApp LUN. For details, see Figure 18 and Table 14.

Special attention isn't required for LUNs attached directly to the child OS if the LUN protocol type of the LUN matches the guest operating system type. For details, see "LUNs" under section 5.2 and follow the instructions in "Selecting the Correct LUN Protocol Type." In addition, see the LUN Multiprotocol Type section in the Data ONTAP Block Access Management Guide for iSCSI or FC or the Data ONTAP Commands Manual Page Reference Document, Volume 1 or Volume 2 on the NetApp Support site.

File System Alignment Process

Aligning the Boot Disk

Virtual disks to be used as boot disks can be formatted with the correct offset at the time of creation by connecting the new virtual disk to a running VM before installing an operating system and then manually setting the partition offset. For Windows guest operating systems, one might consider using an existing Windows Preinstall Environment boot CD or alternative tools such as Bart's PE CD.

To set up the starting offset, follow these steps:

1. Boot the child VM with the Windows Preinstall Environment boot CD.

   Select Start > Run and enter the following command:

```
diskpart
```

2. Type the following into the prompt:

```
select disk 0
```

3. Type the following into the prompt:

```
create partition primary align=32
```

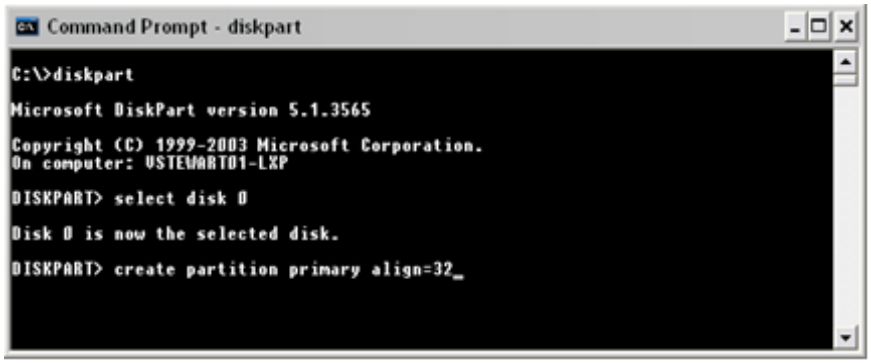

4. Reboot the child VM with the Windows Preinstall Environment boot CD.
5. Install the operating system as normal.

Aligning the DATA Disk

Virtual disks to be used as data disks can be formatted with the correct offset at the time of creation by using Diskpart in the VM.

To align the virtual disk, follow these steps:

1. Boot the child VM with the Windows Preinstall Environment boot CD.
2. Select Start > Run and enter the following command:

```
diskpart
```

3. Determine the appropriate disk to use by typing the following into the prompt:

```
list disk
```

4. Select the correct disk by typing the following into the prompt:

```
select disk [#]
```

5. Type the following into the prompt:

```
create partition primary align=32
```

6. To exit, type the following in the prompt:

```
exit
```

7. Format the data disk as you would normally.

File System Alignment Correction

Detection

When the file systems of virtual disks are aligned for use with NetApp storage systems, the starting partition offset must be divisible by 4,096. For Windows guest operating systems, verifying this value is easy:

1. Run msinfo32 on the guest VM by selecting Start > All Programs > Accessories > System Tools > System Information.
2. Navigate to Components > Storage > Disks and check the value for partition starting offset.

For misaligned VMs, the VM is typically running with a default starting offset value of 32,256, which is not completely divisible by 4,096, and so the partition is not aligned as shown in Figure 19.

Note: For pass-through disks or raw LUNs directly mapped to the VM, 32,256 is reported as the correct starting offset value. This is true because the storage controller compensated for the offset when selecting the correct LUN type.

Figure 19) Using system information to identify the starting partition offset.

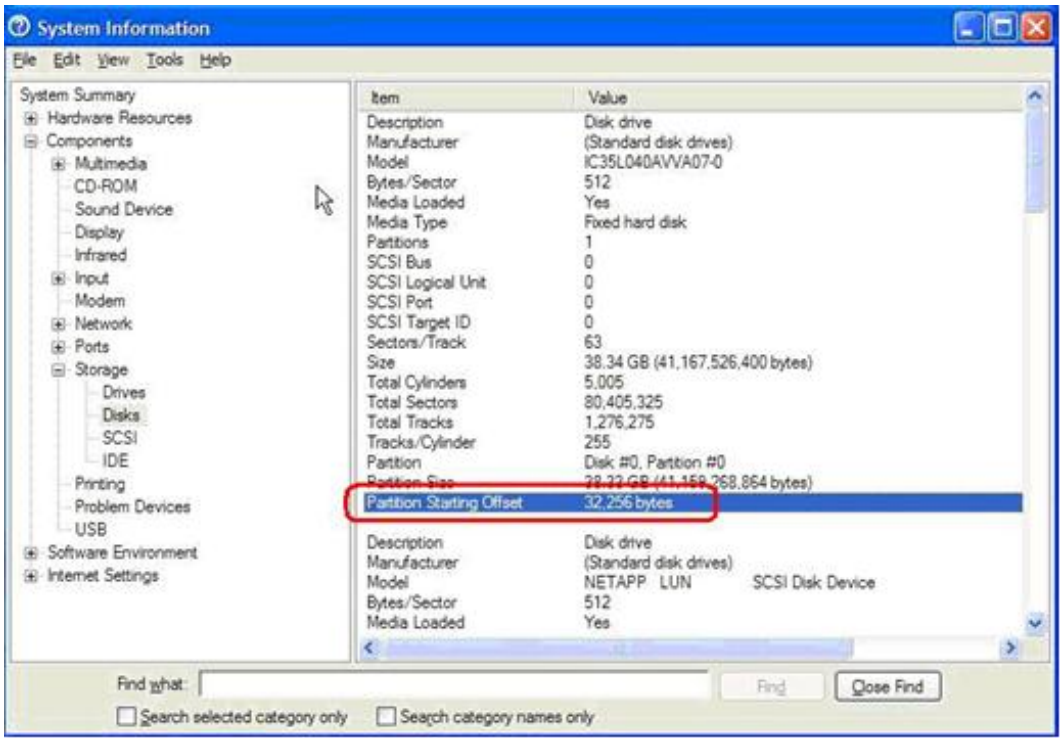

Correction

Correcting the starting offset is best addressed by first correcting the template used to provision new VMs. To do so, follow these steps:

1. Use the procedures described in “File System Alignment Process” earlier in this section to create a new aligned virtual disk.
2. Attach this new aligned virtual disk to the VM.
3. Copy the contents from the existing misaligned virtual disk to the new aligned virtual disk.
4. Detach and destroy the misaligned virtual disk after verifying the contents and integrity of the data on the new aligned virtual disk.

If the misaligned virtual disk is the boot partition, follow these steps:

1. Back up the VM system image.
2. Shut down the VM.
3. Attach the misaligned system image virtual disk to a different VM.
4. Attach a new aligned virtual disk to this VM.
5. Copy the contents of the system image (for example, C: in Windows) virtual disk to the new aligned virtual disk.

Various tools can be used to copy the contents from the misaligned virtual disk to the new aligned virtual disk:

- Windows xcopy
- Norton/Symantec™ Ghost: Norton/Symantec Ghost can be used to back up a full system image on the misaligned virtual disk and then be restored to a previously created, aligned virtual disk file system.

For Microsoft Hyper-V LUNs mapped to the Hyper-V parent partition using the incorrect LUN protocol type but with aligned VHDs, create a new LUN using the correct LUN protocol type and copy the contents (VMs and VHDs) from the misaligned LUN to this new LUN.

For Microsoft Hyper-V LUNs mapped to the Hyper-V parent partition using the incorrect LUN protocol type but with misaligned VHDs, create a new LUN using the correct LUN protocol type and copy the contents (VMs and VHDs) from the misaligned LUN to this new LUN. Next, perform the steps in "File System Alignment Process" earlier in this section to create new aligned VHDs on the new LUN and copy the contents from the existing VHDs to the new VHD. If the VHD is a boot partition, follow the steps described earlier in this section.

For pass-through disks and LUNs directly mapped to the child OS, create a new LUN using the correct LUN protocol type, map the LUN to the VM, and copy the contents from the misaligned LUN to this new aligned LUN.

For more information on disks and storage for Hyper-V, see Planning for Disks and Storage and Configuring Disks and Storage on Microsoft TechNet.

6 Increasing Storage Efficiency and Flexibility

Hyper-V provides an excellent means to increase the hardware utilization of the physical servers. By increasing hardware use, you can reduce the amount of hardware in a data center, lowering the cost of data center operations. In a traditional Hyper-V environment, the process of migrating physical servers to Hyper-V child VMs does not reduce the amount of data stored or the amount of storage provisioned. By default, server virtualization does not have any effect on improving storage utilization (and in many cases might have the opposite effect).

NetApp offers storage virtualization technologies that can further enhance the storage virtualization achieved by these types of disks. These NetApp technologies offer considerable storage savings by providing the capability to thin provision the SAN LUNs and also deduplicate redundant data on them. Both of these technologies are native to NetApp storage systems and don't require any configuration considerations or changes to be implemented in the Hyper-V environment.

6.1 Storage Thin Provisioning

Traditional storage provisioning and the preallocation of storage on disk together are a well-understood method for storage administrators. It is a common practice for server administrators to overprovision storage in order to avoid running out of storage and the associated application downtime required when expanding the provisioned storage. Although no system can be run at 100% storage utilization, there are methods of storage virtualization that allow administrators to address and oversubscribe storage in the same manner as with server resources (such as CPU, memory, networking, and so on). This form of storage virtualization is referred to as "thin provisioning."

Traditional provisioning preallocates storage; thin provisioning provides storage on demand. The value of thin-provisioned storage is that storage is treated as a shared resource pool and is consumed only as each individual VM requires it. This sharing increases the total utilization rate of storage by eliminating the unused but provisioned areas of storage that are associated with traditional storage. The drawback to thin provisioning and oversubscribing storage is that (without the addition of physical storage) if every child VM requires its maximum possible storage at the same time, then there will not be enough storage to satisfy the requests.

NetApp thin provisioning allows LUNs that are presented as physical disks to be provisioned to their total capacity yet consume only as much storage as is required to store the VHD files. In addition, LUNs connected as pass-through disks can also be thin provisioned.

Best Practice

NetApp recommends using thin-provisioned LUNs where possible in the Hyper-V environment for maximum storage efficiency.

However, when enabling NetApp thin provisioning, administrators should also configure storage management policies on the volumes that contain the thin-provisioned LUNs. The use of these policies aids in providing the thin-provisioned LUNs with storage capacity as they require it. The policies include automatic sizing of a volume, automatic NetApp Snapshot copy deletion, and LUN fractional reserve.

Thin Provisioning a LUN

To create a thin-provisioned LUN using NetApp FilerView, follow these steps:

1. Open FilerView (http://<controller IP address>/na_admin).
2. Select LUNs.
3. Select Wizard.
4. In the Wizard window, click Next.
5. For Path, enter the path of the LUN to be created.
6. For Size, enter the size of the LUN to be created and select the measurement to be used for the size of the LUN.
7. For LUN Protocol Type, enter the correct LUN type. For physical disks presented as pass-through disks or direct access by child VM over iSCSI, select the child OS type.
8. For Space Reserved, unselect the checkbox because NetApp recommends not having the LUN as space reserved.
9. For Description, enter an appropriate description for the LUN, if you like.

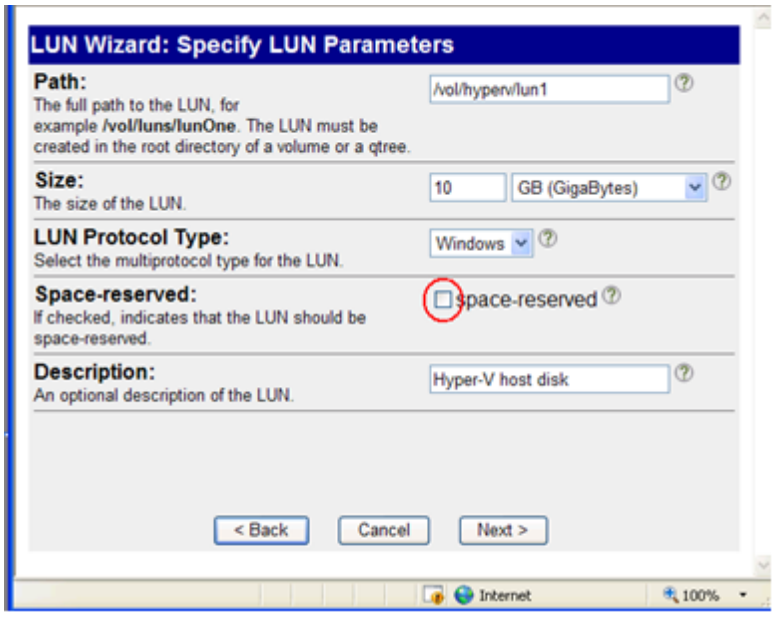

10. Click Finish to complete creation of a thin-provisioned LUN.

Volume Autosize

Volume AutoSize is a policy-based space management feature of Data ONTAP that allows a volume to grow in defined increments up to a predefined limit if the volume is nearly full.

Best Practice

For Hyper-V environments, NetApp recommends setting this value to ON. Doing so requires setting the maximum volume and increment size options.

To enable Volume AutoSize using the NetApp console, follow these steps:

1. Log in to the NetApp console using either SSH, telnet, or a console connection.
2. Type the following into the prompt:

```
set volume autosize policy: vol autosize <vol-name> [-m
<size>[k|m|g|t]] [-i <size>[k|m|g|t]] on.
```

Snapshot AutoDelete

Snapshot AutoDelete is a policy-based space-management feature that automatically deletes the oldest Snapshot copies on a volume when that volume is nearly full.

Best Practice

For Hyper-V environments, NetApp recommends setting the Snapshot AutoDelete value to delete Snapshot copies at 5% of available space. In addition, you should set the volume option to have the system attempt to grow the volume before deleting Snapshot copies.

To enable these options using the NetApp console, follow these steps:

1. Log in to the NetApp console using either SSH, telnet, or a console connection.
2. Type the following into the prompt:

```
set Snapshot autodelete policy: snap autodelete <vol-name> commitment try
trigger volume target_free_space 5 delete_order oldest_first
```

3. Type the following command into the prompt:

```
set volume autodelete policy: vol options <vol-name> try_first volume_grow.
```

LUN Fractional Reserve

LUN fractional reserve is a policy that is required when you use NetApp Snapshot on volumes that contain Hyper-V LUNs. This policy defines the amount of additional space reserved to assure LUN writes if a volume becomes 100% full. LUNs that have space reservation turned off are not affected by the fractional reserve setting.

6.2 NetApp Deduplication

With NetApp deduplication, Hyper-V deployments can eliminate the duplicate data in their environment, enabling greater storage utilization. It can be seamlessly introduced into the Hyper-V environment without having to make any changes to Hyper-V administration, practices, or tasks. Deduplication runs on the NetApp storage system at scheduled intervals and does not consume any CPU cycles on the Hyper-V server. Deduplication can be extremely helpful for scenarios such as fixed-size VHDs, frequent creation and deletion of VHD files on the SAN LUNs, or data in the child VM.

Deduplication is enabled on the NetApp volume, and the amount of data deduplication realized is based on the commonality of the data stored in a deduplication-enabled volume.

Best Practice

For Hyper-V environments, NetApp recommends using deduplication on the NetApp FlexVol volumes containing the LUNs that handle storage for the virtual machines, especially the LUNs provisioned to the Hyper-V servers for VHD storage purposes.

For maximum benefit, consider organizing the VM virtual disks so that VHDs and pass-through disks containing similar child OSs and data reside within the same NetApp FlexVol volume.

Deduplication Considerations

Enabling deduplication when provisioning LUNs results in storage savings. However, the default behavior of a LUN is to reserve an amount of storage equal to the provisioned LUN. This design means that although the storage array reduces the amount of capacity consumed, any gains made with deduplication are for the most part unrecognizable because the space reserved for LUNs is not reduced.

To recognize the storage savings of deduplication with LUNs, you must enable NetApp LUN thin provisioning. For details, see section 6.1, “Storage Thin Provisioning.” In addition, although deduplication reduces the amount of consumed storage, the Hyper-V administrative team does not see this benefit directly, because their view of the storage is at a LUN layer, and LUNs always represent their provisioned capacity, whether they are traditional or thin provisioned.

A deduplicated flexible volume can be replicated to a secondary storage system (destination) using NetApp SnapMirror technology. The flexible volume at the secondary site inherits all the deduplication attributes and storage savings through SnapMirror, thereby reducing cost. Also, only unique blocks are transferred to the secondary site, so deduplication reduces network bandwidth usage as well.

Enabling NetApp Deduplication

To enable and initialize NetApp deduplication on a flexible volume, follow these steps:

1. Log in to the NetApp console using either SSH, telnet, or a console connection.
2. To enable deduplication, type the following into the prompt:

```
sis on <volume path>
```

3. To start processing existing data, type the following into the prompt:

```
sis start –s <volume path>
```

4. To monitor the status of NetApp deduplication operation, type the following into the prompt:

```
sis status
```

For deduplication best practices, including scheduling and performance considerations, see TR-3505: NetApp Deduplication for FAS: Deployment and Implementation Guide.

6.3 NetApp FlexClone Technology

NetApp FlexClone® technology creates true cloned volumes, which are instantly replicated datasets, files, LUNs, and volumes without additional storage space utilized at the time of creation. FlexClone volumes are writable point-in-time copies generated from the Snapshot copy of a FlexVol volume. A FlexClone volume has all of the features of a FlexVol volume, including growing, shrinking, and being the base for a Snapshot copy or even another FlexClone volume.

FlexClone volumes deployed in a Hyper-V virtualized environment offer significant savings in dollars, space, and energy. Additionally, a FlexClone volume or file has performance identical to that of any other FlexVol volume or individual file.

Best Practice

For Hyper-V environments, NetApp recommends using FlexClone volumes for a variety of tasks, including constructing test and development environments, rapid VM provisioning for virtual desktop infrastructures (VDI) scenarios, and more.

FlexClone Concepts

Deploying FlexClone volumes offers great results in circumstances where the data growth is incremental and where the information is distributed in changeable form without hampering the integrity of the original data.

The FlexClone volume creation is instantaneous and doesn't affect the accessibility of the parent FlexVol volume. These volumes are space-efficient; only as the parent volume and the clone diverge due to variations in the clone's data blocks are those new blocks and their block-map pointers written to disk, so the volume/LUN accumulates only the changed blocks. The setup can be coupled with a NetApp replication solution such as SnapMirror to create incrementally propagated data on the secondary storage system.

FlexClone Volume Creation

Create a FlexClone Volume from the Console

FlexClone volumes can be created using either the shell console or FilerView. To create a FlexClone volume of a FlexVol volume, follow these steps:

1. Log in to the NetApp console using either SSH, telnet, or a console connection.
2. To create a FlexClone volume, type the following into the prompt:

```
vol clone create FlexCloneName –b ParentFlexVolName
```

Create a FlexClone Volume from NetApp SnapDrive

FlexClone volumes can be created using either the shell console or FilerView. To create a FlexClone volume of a FlexVol volume, follow these steps:

1. Access the SnapDrive console.
2. Click the Connect Disk option.
3. Click Next on the Connect Disk Welcome wizard.
4. Provide the details (IP address or hostname) of the storage system and establish the connection.
5. Select the volume and then drop down through the Snapshot container to select the desired Snapshot copy, and finally, the desired LUN as indicated in the following screenshot.

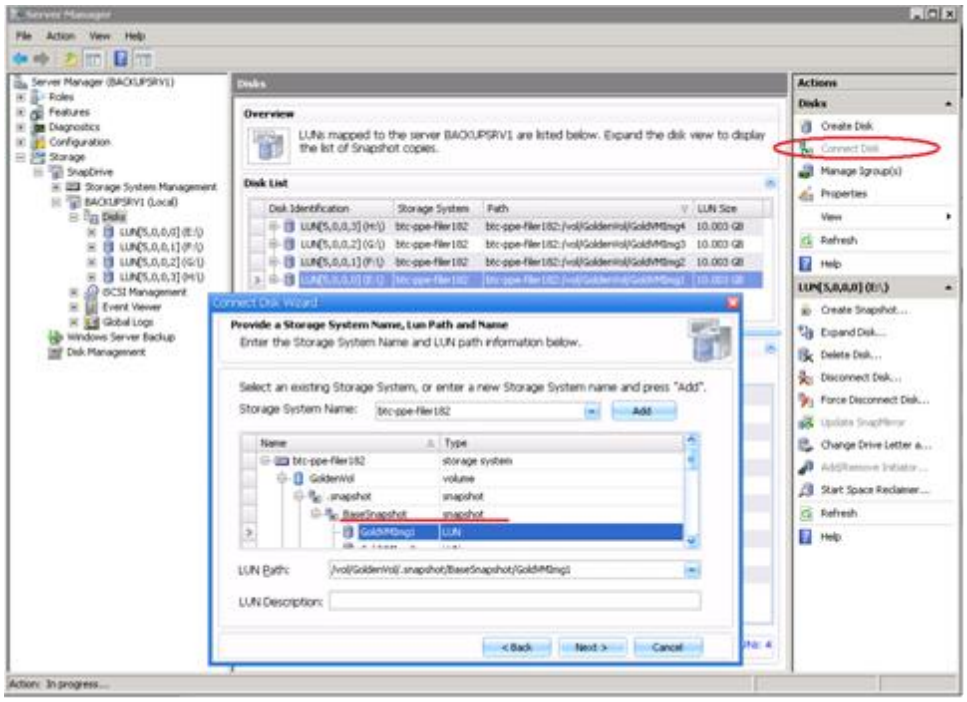

6. Click Next to continue.
7. Allow the LUN type to be Dedicated.
8. Click Next to continue.
9. Specify a drive letter or volume mount point to connect the LUN.
10. Click Next to continue.
11. From the list of initiators, select the checkbox of the WWPN or IQN of the initiator for which the LUN is intended.
12. Click Next to continue.
13. Allow the Initiator Group Management to be set to Automatic.
14. Click Next to continue.
15. Verify the settings of the created disk.
16. Click Finish to complete the task.

6.4 NetApp Snapshot Copies

A NetApp Snapshot copy is a locally retained, frozen, space-efficient read-only view of the volume or an aggregate. It facilitates with improved stability, scalability, recoverability, and performance that are better than those provided by any other storage Snapshot technologies.

Snapshot copies facilitate frequent, low-impact, user-recoverable online backups of files, directory hierarchies, LUNs, and application data. They provide a secure and simple method of data restores through which users can directly access the Snapshot copies and recover from accidental file deletions, data corruptions, or modifications. The SnapManager suite of products, which is available for various enterprise applications, uses the features of Snapshot copies and delivers an enterprise-class data protection solution.

Snapshot Concepts

Snapshot is an industry-known technology that involves the facility to create a temporary point-in-time recovery copy. Microsoft Hyper-V offers the capability of creating snapshots of the VMs that it hosts. NetApp Snapshot technology tightly integrates with virtual infrastructure deployed using Hyper-V and complements the server virtualization objective. It provides a crash-consistent, point-in-time recovery copy of the VM image that is useful in case of an entire VM restore, VM cloning, site replication, or data recovery.

The Hyper-V snapshot creation initiated from Hyper-V Manager, from SCVMM, or through Windows PowerShell uses the hardware resources from the Hyper-V server, whereas the NetApp Snapshot copy creation offloads the task execution to the storage system so that the host resources don't take any performance hits. Increasing the number of Hyper-V snapshots affects the performance of the host system, whereas the NetApp storage system can handle up to 255 Snapshot copies per volume with no performance degradation and minimal consumption of storage space. NetApp offers SnapManager for Hyper-V (SMHV) to achieve this.

The Snapshot copies created by SMHV can be backed up to tape or any secondary storage system and replicated to another facility with NetApp SnapMirror or SnapVault. VMs can be restored almost instantly, individual files can be quickly and easily recovered, and clones can be instantly provisioned for test and development environments. Refer to the later chapters of this document (Chapters 12 and following) for a detailed explanation of SMHV.

Best Practice

For Hyper-V environments, NetApp recommends using SnapManager for Hyper-V for a variety of tasks, including backup and recovery of VMs, disaster recovery of VMs, and more.

7 Virtual Machine Provisioning

Virtual infrastructure solutions, such as Microsoft Hyper-V, empower IT organizations to rapidly deploy virtual machines in all phases: development, test, and production. The tasks involved to deploy virtual machines usually generate many physical copies of virtual machine images, which demands more storage resources to maintain the many virtual machine instances and management resources to execute the many manual steps required to deploy these virtual machines individually.

Integration of Microsoft virtual environments with NetApp storage technology can solve these challenges by helping organizations reduce the efforts spent deploying individual virtual machines and reduce the amount of storage required to support the deployment of individual virtual machines, which helps reduce costs associated with space, power, and cooling. Use of NetApp Snapshot and FlexClone technology, along with NetApp deduplication, can support the rapid deployment of tens, hundreds, and thousands of virtual machines in minutes while minimizing the total storage supporting such a deployment by 50% or more when compared to a baseline of traditional storage.

7.1 Provisioning Concepts

NetApp Snapshot and FlexClone

The traditional VM provisioning process involves tedious and time-consuming tasks such as provisioning storage, installing the OS environment, patching it up with required service packs and applications, and rolling it out to the end user. NetApp storage, with its Snapshot and FlexClone technologies, facilitates an instantaneous zero–space consuming writable copy of the flexible volumes. These FlexClone volumes can be provisioned within a matter of minutes. They contain LUNs with virtual hard drives (VHDs) that can be connected to and used as individual OS instances for virtual machines.

NetApp Deduplication for Space Savings

Creating multiple copies of the golden virtual machine images specific to user requirements can result in higher space utilization. Deduplication technology operates at the block level, eliminating data blocks with identical content, and maintaining a single copy of the dataset, thereby achieving a higher level of space savings. Deduplication can be implemented in a production environment with minimal effect on actual storage performance and in some environments can even increase actual storage performance.

7.2 Virtual Machine Provisioning Process

Figure 20 shows the process flow required to provision Hyper-V VMs using NetApp cloning techniques.

Figure 20) Process flow to provision Hyper-V VMs using NetApp cloning techniques.

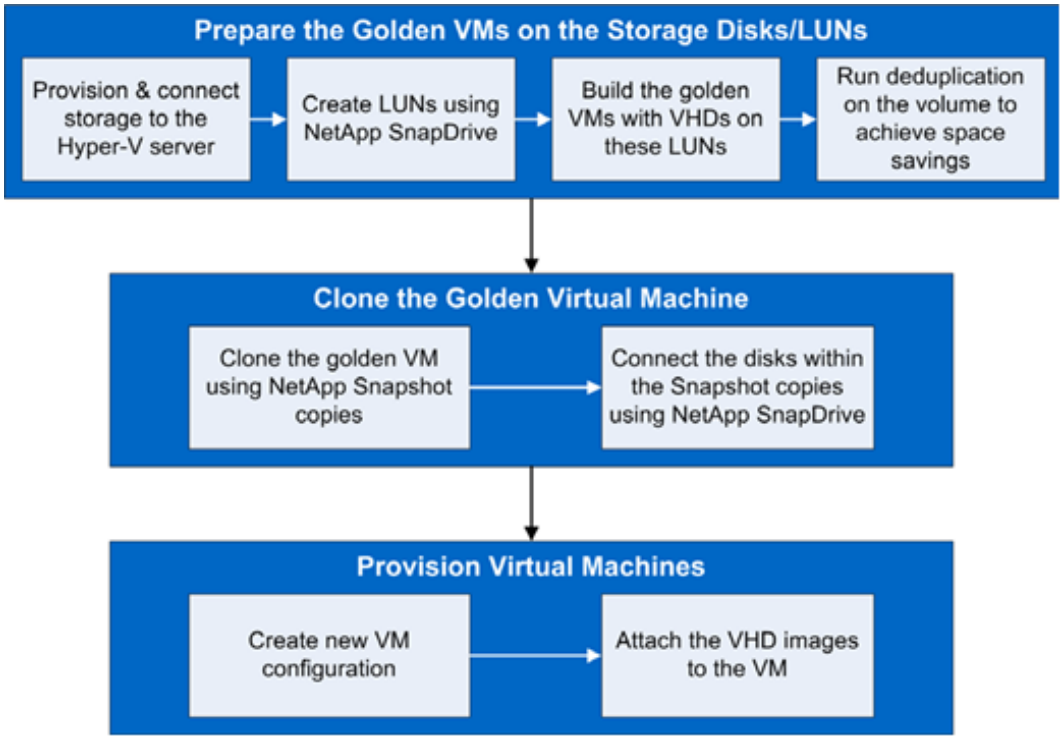

Preparing the Golden Virtual Machine

In cases where the infrastructure requires multiple copies of the OS instance, it is a repetitive and time-consuming task to perform installation. Server virtualization offers an efficient method to reduce this task and allows attaching an existing image as an OS disk. Using this facility, we can perform a one-time installation of the child OS and designate it as a golden image. The golden image can then be updated at any time with required service packs and applications as needed. This lets administrators provision desktops and servers to users in matter of minutes.

To provision storage for the Hyper-V Server, follow these steps:

1. Create the aggregate.

   Follow the NetApp best practice recommended settings for new aggregates described in section 5.2 under “Aggregates.”
2. Create the golden volume.

Create a new volume within the aggregate created in the previous step and follow the NetApp best practice recommended settings for new flexible volumes described in section 5.2 under Flexible Volumes."

To connect storage to the Hyper-V Server:

1. Connect to the NetApp storage using NetApp SnapDrive.

NetApp SnapDrive for Windows can be used on the Windows 2008 R2 Server to manage the NetApp storage system. For details on best practices for configuration and use of NetApp SnapDrive for Windows, see "NetApp SnapDrive for Windows" in section 5.1. For details on the installation of NetApp SnapDrive for Windows, refer to the "NetApp SnapDrive for Windows" section of TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide, and follow the instructions in the "Installation" section.

If using iSCSI connectivity between the Hyper-V server and NetApp storage, be sure to establish an iSCSI session before completing the next steps.

2. Create a LUN using NetApp SnapDrive in the golden volume.

   Create a new LUN within the flexible volume created in the previous step and follow the NetApp best practice recommended settings for new LUNs described in "LUNs" in section 5.2, specifically the subsection titled "Process to Create a LUN."

To build the golden virtual machine:

1. Create a VM using Hyper-V Manager or SCVMM.

   For details, see the "Virtual Machine Provisioning" section of TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide and follow the instructions in "Provisioning Using Hyper-V Manager or Provisioning Using SCVMM 2008."

   Create a virtual machine with a fixed VHD (virtual hard disk) of an appropriate size for the expected operating system on the LUN created in the previous step.

2. Install the child operating system.

   For details, see the "Install Operating System" section of TR-3701: NetApp Implementation Guide for Microsoft Virtualization and follow the instructions.

3. After completing the installation process, NetApp recommends installing the Hyper-V Integration Services, also known as integration components (IC).

   The ICs are installed for the time synchronization, heartbeat, shutdown, key/value pair exchange, and Volume Shadow Copy Service purposes. For details, see the "Install Hyper-V Integrated Services" section of TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide.

4. Install all applications and necessary OS and application updates.

   The virtual machine image is a golden copy that will be further used to provision identical VMs. To avoid repetitive application installations, it is a good idea to install all OS-related patches and required applications such as service packs, antivirus applications, office automation software, and so on.

5. Configure the child OS and shut down the VM.

   Configure VMs provisioned from a golden image using the Microsoft System Preparation (SysPrep) Tool before pushing it into production. It is a process to generate a secure ID (SID) for the operating system instance so that it remains unique. Refer to Microsoft KB 302577 for the detailed instructions on its usage.

6. Enable NetApp deduplication on the volume.

   Multiple LUNs can be created in a single FlexVol volume and copies of VHDs can be stored on these LUNs, which would be attached to the Hyper-V server as physical disks. Each Hyper-V virtual machine might have the same or different set of applications installed within the operating system environment (OSE) as needed. Space savings can be achieved with the NetApp deduplication capabilities. For details, see "Enabling NetApp Deduplication" in section 6.2 and follow the instructions.

Before enabling deduplication on the flexible volume, make sure that best practices have been followed to disable space reservation on all LUNs within the flexible volume by unselecting Space Reserved in the LUN properties. For details, see “LUNs” in section 5.2, specifically the subsection titled “Process to Create a LUN.”

To clone the golden VM:

1. Clone the golden VM using NetApp Snapshot technology.

   NetApp FlexVol volumes can be cloned with zero space consumed by creating NetApp Snapshot copies. An individual connection can be established to the LUNs existing on these clones to attach them as separate physical disks to the Hyper-V server.

   Within NetApp SnapDrive, select the physical disk on which the golden VM image resides and create a NetApp Snapshot copy of this physical disk.

2. Connect the disks in the NetApp Snapshot copy using NetApp SnapDrive.

   After the NetApp Snapshot copy has been successfully created, we can use SnapDrive to connect to the individual LUNs within the Snapshot copy as individual disks. We will create FlexClone volumes of this Snapshot copy. For details, see “FlexClone Volume Creation” in section 6.3.

To provision the VMs:

1. Create a new VM.

   The Hyper-V Server Manager or SCVMM can be used to create the desired number of virtual machines. For details, see the “Virtual Machine Provisioning” section of TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide and follow the instructions in “Provisioning Using Hyper-V Manager or Provisioning Using SCVMM 2008.”

   When creating the new VMs, be sure to use the option `Attach a Virtual Disk Later` so that you will have blank VM configurations ready for the cloned VHD.

2. Attach the golden VHD images existing on the FlexClone volume.

   Within the settings for each virtual machine (right-click the virtual machine), under the IDE Controller 0, add a hard disk and browse to the location of the VHD stored on the NetApp FlexClone volume. It is a good idea to rename the VHD to match the virtual machine name or disk (for example: `[VMname]_Vol0`) before connecting it to the VM. The VM can then be powered on as usual and given a unique configuration (host name, IP address, and so on).

8 Backup and Recovery

Backup and recovery is the most critical component of the data protection plan. If data is changed unexpectedly, a system is compromised, or a site is lost, then backup comes in handy to protect and recover business information assets.

NetApp backup and recovery solutions equip users to increase the reliability of data protection while minimizing the management overhead and cost involved. These solutions fit into any strategy, enabling users to meet the service-level requirements.

8.1 Storage Considerations for Virtual Machine Backup and Recovery

Virtual Machine Configuration

Datasets growing in the virtualized environment lead to a huge amount of transient and temporary data on the virtual disks. It is required to separate this data while deploying with NetApp Snapshot and SnapMirror solutions. Since the Snapshot copy holds onto the storage blocks that are not used anymore, the transient and temporary data can consume a large amount of space in a short time period. Hence, it pays off to separate the valuable data from the transient data, which accounts for the data sent during replication updates or backups done for data protection.

The virtual machine page files as well as the user and system temp directories must be created on separate virtual disks residing on a unique datastore dedicated to the respective directory type.

Virtual Machine Data Layout

Virtual machines created on the Hyper-V servers account for the storage requirements of their configuration files `<vmguid>.bin` and `<vmguid>.vsv`.

The .bin file is used to allocate enough space to store the contents of memory of the virtual machine when the VM needs to be saved, and the file size will be the amount of memory defined for the child partition. The .vsv file is for storing saved state of devices. When booting a virtual machine, 20MB is reserved for the .vsv. When placing the virtual machine into a saved state, this file might shrink or grow to 50MB depending on the number of devices configured for the virtual machine. By default, these are stored in the VM configuration directory and should be placed in a separate datastore and a different volume.

As seen from the following block diagram, the virtual machine is created to have its configuration files on a disk within volume `FlexVol2`, which has Snapshot copies disabled. This volume additionally has a disk that is presented as `Drive D` to the virtual machine, which can be used to store files and system temp directories. The setup makes sure that the transient and temp data are segregated into a separate volume that doesn't need backup.

The other volume, `FlexVol2`, has Snapshot copies enabled. This volume contains the LUN presented as `Drive C` to the virtual machine and can be actively backed up using the NetApp Snapshot features.

Figure 21 illustrates the guest VM file system aligned with the storage array blocks.

Figure 21) Guest VM file system aligned with the storage array blocks.

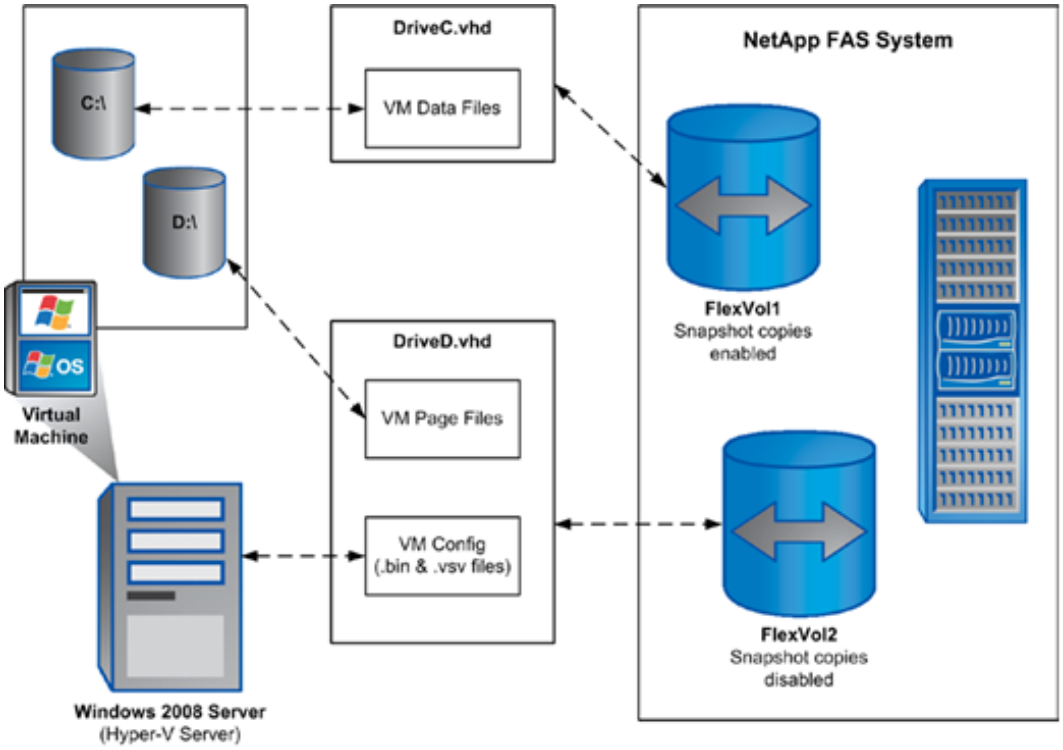

8.2 Backup Using NetApp SnapManager for Hyper-V

Data protection plans for a virtualized environment get more critical and inevitable as the consolidation brings all crucial data into one place, and any form of failure results in a massive impact on the business applications.

Backup tasks running in the server virtualized infrastructure are often resource-intensive (CPU, memory, disk I/O, network) and result in bottlenecks that adversely affect the performance of the other business-critical applications sharing the same environment. Backup schedules must be closely coordinated with those applications running on the available resource. NetApp offers SnapManager for Hyper-V (SMHV) to enable application-consistent and VM-consistent backup and restore. It also facilitates disaster recovery by leveraging SnapMirror technology and Windows PowerShell cmdlets. SnapManager for Hyper-V is discussed in detail later in this report (Chapters 12 and following).

9 Disaster Recovery and High Availability

Business operations are heavily dependent on the information systems and the related IT infrastructure. A minor application outage might cause a significant impact, and effect of a data loss is even more critical. Various metrics are commonly used in designing a business continuity plan.

Two of the most frequently used metrics are recovery point objective (RPO) and recovery time objective (RTO). RPO (measured in minutes and hours) describes how far the recovered data is out of sync with the production data at the time of disaster. RTO (measured in minutes) describes how quickly the operations can be restored.

Several approaches have been followed to increase data availability and business continuity against disaster occurring at the hardware and the software level, or even site failures. Primarily, backup methods provide a way to recover from the data loss from an archived medium that offers a high-level data protection method. Redundant hardware setups can provide a second level of protection to mitigate any damage caused by the hardware failures.

Data mirroring is another mechanism to maintain data availability and minimize downtime. The NetApp SnapMirror solution empowers IT infrastructures with a fast, flexible data replication mechanism over Ethernet and Fibre Channel networks. It is a key component to be considered while designing and deploying enterprise data protection plans. SnapMirror can function as an efficient data replication solution since it can take advantage of underlying NetApp technologies such as Snapshot copies, FlexClone volumes, deduplication, and so on.

Disaster recovery being its primary objective, SnapMirror can also assist other critical application areas such as DR testing, application testing, load sharing, remote tape archiving, and remote data access.

9.1 Business Continuance Concepts

Disaster is inevitable for any IT infrastructure, which makes it more critical for environments to be consolidated using server virtualization because the consolidation would bring in additional complexity by sharing reduced physical hardware resources for the applications and the business-critical data that is running within the OS instances. The infrastructure must be designed with special attention to take care of the challenges that arise in a virtualized environment.

Some of the challenges are as minor as the following:

- There is little or no time to schedule a downtime window to perform cold backup on virtual machines.
- Hot backup of virtual machines results in inconsistent backup copies, which are useless during recovery.
- Infrastructure with various OS instances presents difficulties in identifying the consistent state of the backup.

- Replicating data over LAN/WAN might consume more than the available resources.
- Planning for identical resources at the DR site results in an increase of TCO and unused infrastructure.

NetApp offers solutions that complement the server virtualization solutions and help to mitigate these challenges. Solutions such as NetApp Snapshot copies, FlexClone volumes, and deduplication enable an architect to design a complete data protection solution and also by make use of the available resources efficiently.

NetApp Snapshot Copies

NetApp Snapshot copies are the base for the family of SnapManager data protection solutions offered for enterprise applications. Hence, each SnapManager family of products inherits the unique advantages of NetApp Snapshot technology. The data replication solution offered by NetApp functions on the basis of mirroring the point-in-time recovery copies. The subsequent Snapshot copies are mirrored with the blocks that were added or changed since the previous Snapshot copies. This incremental behavior limits the associated storage consumption by offering effective resource utilization in terms of storage space and network bandwidth.

NetApp FlexClone Volumes

NetApp FlexClone volumes offer a flexible data management solution. A FlexClone volume can be created instantaneously from any FlexVol volume or Snapshot copy with which it can transform a read-only Snapshot copy into a read-write copy. These copies are thin provisioned, which implies that the space usage is only increments of the delta data. The FlexClone capability of NetApp storage helps to test the behavior of the HA and the DR setup. It is possible to create a read-write copy of the Snapshot copy available at the DR site and test its availability in case of a real disaster.

NetApp Deduplication

NetApp deduplication offers an excellent solution to eliminate the data saved on the duplicate blocks and consolidates them into a single block. This results in an efficient method of storage space utilization. This solution further transforms to be effective for the HA and DR solutions designed with the NetApp replication technologies since the overall amount of data that needs to be mirrored to the destination storage system is reduced.

9.2 NetApp SnapMirror

NetApp SnapMirror software is a simple, flexible, cost-effective disaster recovery and data distribution solution deployed for more of the enterprise application infrastructure. Data replication happens across the LAN or WAN, offering high availability and faster disaster recovery for the business-critical applications. Continuous data mirroring and mirror updates happening across multiple NetApp storage systems would facilitate the mirrored data for multiple purposes. Businesses that are spread out geographically can take advantage of SnapMirror and make local copies of mirrored data available to all locations, highly enhancing efficiency and productivity.

Simplified Deployment and Administration Procedure

The SnapMirror solution can be deployed within a matter of minutes, and administration is easy. It offers built-in SNMP support enabling easy integration with the SNMP framework. Because of the simplified administration, there are few chances of operator error during the recovery process. SnapMirror integrates with the SnapManager suite of products to make sure of an application-consistent replica. Various modes of SnapMirror such as Sync, Semi-Sync and Async are available from a single license with the flexibility to choose the level of RPO (0–24 hours) per business needs.

Efficient Resource Utilization

Data replication can happen across FC storage to less expensive SATA storage, which means reduced bandwidth requirements and costs by leveraging third-party WAN acceleration technologies along with NetApp Snapshot technology to send only the changed data blocks.

Mirrored data available on the remote sites can be put to active business use such as running backups, app testing and QA, and staging. SnapMirror has the capability to automatically perform checkpoints during data transfers to eliminate the need for full transfers while recovering from a broken mirror or loss of synchronization by performing intelligent resynchronization.

Enhanced Storage Availability

SnapMirror complements the virtualized infrastructure and helps maintain higher availability for the VMs running within Hyper-V. SnapMirror can be integrated with the available host-level failover cluster utilities and offers a robust disaster-recovery solution.

SnapMirror for Hyper-V Infrastructure

As described in “Virtual Machine Configuration” in section 8.1, virtual machines created on Hyper-V servers account for the temporary and transient data that gets created during its operation. Hence, hosting this data on a separate LUN and excluding it from the mirroring process would result in an efficient replication setup.

Hyper-V offers an option to export the VM configuration to a directory location. This option can be used to save the VM configuration and can be imported during the VM restore. The drive to which the VM configuration is exported can also be replicated to the recovery site so that this data can be used during the restore process. Figure 22 represents a high-level view of the solution.

Figure 22) Showing Hyper-V solution using SnapMirror for disaster recovery.

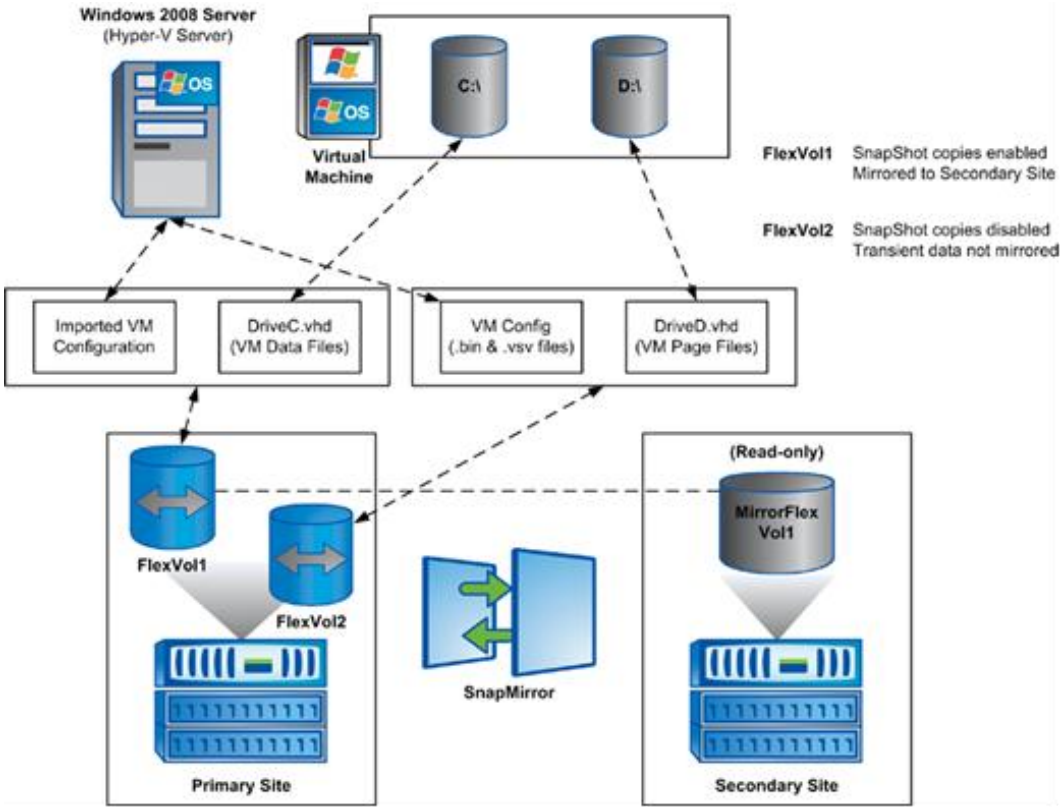

Key Technical Concepts

NetApp SnapMirror offers a replicating solution that operates in synchronous and asynchronous modes.

- In synchronous mode, the data updates are sent from the source to the destination as they occur. Synchronous mode is used for critical business applications that need near 100% availability and the lowest RPO and RTO, but the distance between the source and destination controllers is limited to 100 kilometers.
- In asynchronous mode, the replication of the data between the source and destination is independent of the data changes in the source controller. The changes are replicated on a scheduled predetermined time to the destination storage system. There are no distance limitations in this mode, with RPO and RTOs to minutes and hours, but there is some exposure to data loss.

Synchronous Replication

Synchronous replications can be implemented in data centers that have strict uptime requirements. This mode of replication sends updates from the primary to the remote site as they occur instead of on a predetermined time schedule. This is achieved by replicating every data write to the remote location and not acknowledging to the host that the write occurred until the remote system confirms that the data was written. This solution provides the least data loss, but there is a limit of 50 to 100Km before latency becomes too great because the host application must wait for an acknowledgement from the remote NetApp devices. There is data protection in case of an entire site failure. Since the replication is synchronous, it can have significant performance impact and is not necessary for all applications.

Synchronous replication begins with a one-time baseline transfer in which the entire dataset is mirrored from the primary to the remote site. During this action, the SnapMirror status is indicated as `Transferring`. After the successful completion of the baseline transfer, SnapMirror transitions into synchronous mode, and the status is indicated as `In-Sync`.

Refer to TR-3326: SnapMirror Sync and SnapMirror Semi-Sync Overview and Design Considerations (for the detailed information on the procedure that SnapMirror uses for data replication.

Asynchronous Replication

SnapMirror in asynchronous mode can operate at both volume and qtree levels. It can perform incremental, block-based replication as frequently as once per minute or as infrequently as days apart. There is no distance limitation, and it is frequently used to replicate long distances to protect against regional disasters.

Before the incremental updates can occur, the asynchronous mode begins with a one-time baseline transfer in which the entire dataset is replicated from primary to the remote site. After the successful completion of the baseline transfer scheduled or manually triggered SnapMirror updates, transfer only the changed data blocks. Because only the blocks that have changed between each replication are sent, there is a minimal impact on the write throughput and latency.

Refer to TR-3446: SnapMirror Async Overview and Best Practices Guide for detailed information on the procedure that SnapMirror uses for data replication.

Storage Replication Using NetApp SnapMirror

As mentioned, SnapMirror operates in either synchronous or asynchronous mode. To understand the process involved in configuration of these modes, let’s consider the scenarios of having SnapMirror set up in intrasite and intersite. Let’s set up SnapMirror in synchronous mode for an intersite scenario and configure SnapMirror in asynchronous mode for an intrasite scenario.

Figure 23) Solution for intrasite replication using synchronous SnapMirror.

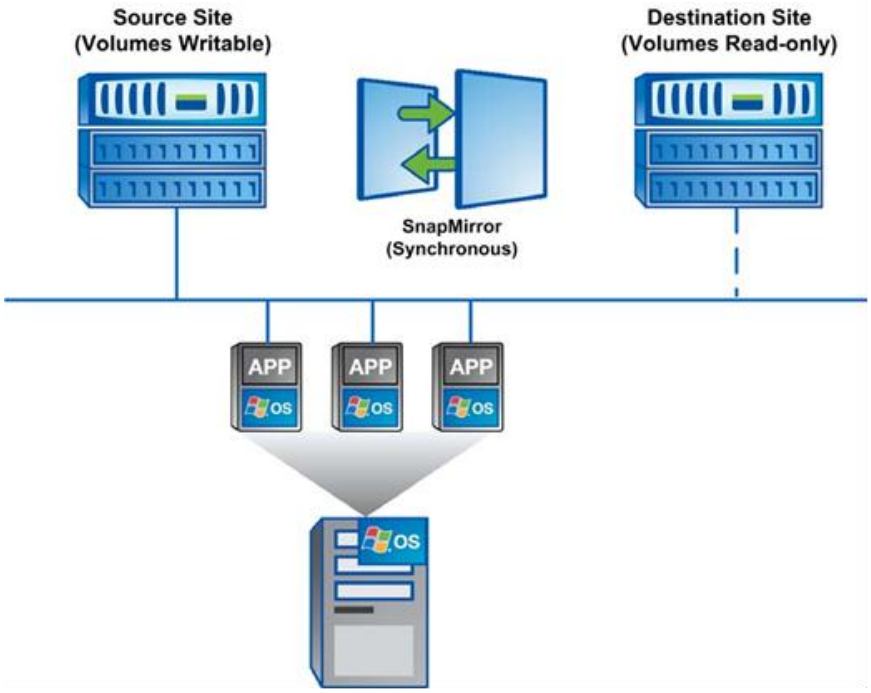

Figure 23 represents a high-level view of the components in an intrasite replication. The Hyper-V server host runs multiple VM instances with their disks hosted on the volumes residing on the source site storage system, and these volumes are writable. The source storage system has established the synchronous SnapMirror relationship with another storage system residing on the same site. The mirrored volumes on the destination site are read only and become available for use only after the SnapMirror relationship is broken between them.

9.3 Configuring NetApp SnapMirror Replication for VMS Between NetApp Storage Systems

To configure the NetApp SnapMirror replication for the virtual machines, you must identify the volumes on which the VHDs of the VMs are hosted. These volumes must be included in the configuration setup.

Use the following procedure to set up a synchronous SnapMirror relationship between the source and the destination storage systems. All of the following commands must be executed from the console of the storage system, which can be accessed using telnet or SSH applications.

1. Be sure to have SnapMirror licenses on both the source and the destination storage systems. In case there is no license, procure, and add them using the following command:

```
sourceFiler> license add snapmirror_license_code destnFiler> license add snapmirror_license_code
```

2. From the source storage system console, specify the hosts that are allowed to use SnapMirror. Use the following command:

```
sourceFiler> options snapmirror.access host=destnFiler
```

3. For a synchronous SnapMirror configuration, add the following entries to the file/etc/snapmirror.conf on the destination storage system:

```
sourceFiler:vhdVolume0  destnFiler:vhdVolume1   -     sync
```

This configuration means that `vhdVolume0` on the source storage system (`sourceFiler`) and the `vhdVolume1` on the destination storage system (`destnFiler`) have a synchronous SnapMirror relationship.

4. Make sure that the size of volume on the destination storage system (vhdVolume1) is equal to or greater than vhdVolume0 and is in restricted mode. Use the following command to switch the destination volume to restricted mode:

```
destnFiler> vol restrict vhdVolume1
```

5. Enable SnapMirror on both the source and the destination storage system using the following command:

```
sourceFiler> snapmirror on destnFiler> snapmirror on
```

6. To initialize the SnapMirror mirroring process, use the following command:

```
destnFiler> snapmirror initialize –S sourceFiler:vhdVolume0
destnFiler:vhdVolume1
```

This command creates an initial (baseline) complete copy of the source volume (`vhdVolume0`) and initializes the mirroring process.

The storage system can function normally while the mirroring process is running in the background.

7. Use the following command to monitor the mirroring status:

```
destnFiler> snapmirror status
```

The output of the command would indicate that the status is `Transferring` while the baseline copy is in progress. It would show the status as `Mirrored` after the completion of the baseline copy.

With this configuration setup, the synchronous SnapMirror relationship has been set between the source and destination volumes of the storage system.

9.4 Disaster Recovery Using NetApp SnapMirror

This section discusses the process of configuring storage replication intersite (between two different physically located sites) using NetApp SnapMirror to protect data from site failure.

Figure 24) Solution for intrasite replication using asynchronous SnapMirror.

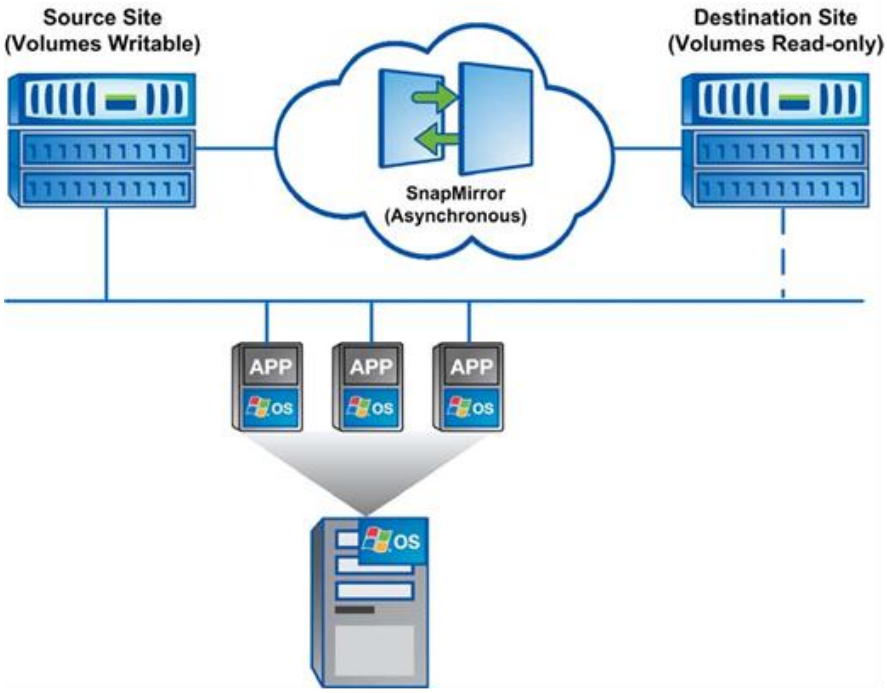

Figure 24 represents a high-level view of the components in an intersite replication. The Hyper-V server host runs multiple VM instances with their disks hosted on the volumes residing on the source site storage system, and these volumes are writable. The source storage system has established the asynchronous SnapMirror relationship with another storage system residing on the same site. The mirrored volumes on the destination site are read only; they become available for use only after the SnapMirror relationship is broken between them.

9.5 Configuring NetApp Snapshot Backups for VMS on the Primary NetApp Storage System

Launch the following command from the command prompt of the Hyper-V server. Use the drive letter on which the virtual machine disks reside.

```
C:\> sdcli snap create –s <snapshot-name> -D <Drive> -x –u yes
```

- `-x`: A Snapshot copy is created only for the LUNs specified in the Drive list.
- `-u (yes / no)`: Setting this option to yes initiates a SnapMirror update.

For example: `sdcli snap create –s VM01-Snapshot1 –D F –x –u yes`

9.6 Configuring NetApp SnapMirror Replication for VMS Between NetApp Storage Systems

The intersite replication scenario should include the procedure to configure asynchronous SnapMirror replication on the volumes that host the vhd files of the virtual machines.

Use the following procedure to set up an asynchronous SnapMirror replication between the source and the destination storage system.

1. Edit the configuration file /etc/snapmirror.conf on the destination storage system to include the following line:

```
sourceFiler:vhdVolume0     destnFiler:vhdVolume1 kbs=2000,restart=always
15 * * 1,2,3,4,5
```

This entry specifies that the Snapshot copy mirroring occurs from `sourceFiler - vhdVolume0` to `destnFiler - vhdVolume1` at a maximum rate of 2,000 kilobits per second, 15 minutes past every hour, Monday through Friday.

2. To initialize the SnapMirror mirroring process, use the following command:

```
destnFiler> snapmirror initialize –S sourceFiler:vhdVolume0
destnFiler:vhdVolume1
```

This command creates an initial (baseline) complete copy of the source volume (`vhdVolume0`) and initializes the mirroring process.

The storage system can function normally while the mirroring process is running in the background.

3. Use the following command to monitor the mirroring status:

```
destnFiler> snapmirror status
```

The output of the command indicates that the status is `Transferring` while the baseline copy is in progress. It would indicate the status as `Mirrored` after the completion of the baseline copy.

9.7 Restoring Service for VMS Between NetApp Storage Systems

In case of a disaster such as accidental file deletion, data loss, or site failure, data available at the destination site can be retrieved.

With the disaster occurring on the source site, the volumes and the LUNs go disconnected, and all virtual machines go into a critical state. The data can be recovered from the destination site by connecting to the volumes available at the destination storage system.

Use the following steps to connect the volumes available at the destination site.

1. Launch the following command from the command prompt of the Hyper-V server:

```
C:\> sdcli disk connect -p $LUNPath -d $VolMountPt -I $fqdn $iqn - dtype
dedicated
```

   Where:

   - $LUNPath indicates the path of the LUN on which the VHDs reside.
   - $VolMountPt indicates the drive/volume mount point on which the LUN is mounted on the host.
   - $fqdn indicates the fully qualified domain name (host name) of the server.
   - $iqn indicates the IQN node name iSCSI/WWPN (FC) of the adapters on the host system.

2. After successful connection to the corresponding LUN, a reset can be done on the virtual machines to bring them back into production.

Information

Orchestration software such as Microsoft System Center Opalis can be used to initiate a disaster recovery by leveraging SnapMirror. Refer to http://communities.netapp.com/docs/DOC-8153 for more details.

10 Monitoring and Managing

Storage monitoring and management are very critical to the successful operations of a Hyper-V environment. NetApp offers the appropriate tools to monitor the health of storage systems, provide alerts, generate reports, and manage storage growth. This section discusses the options available to address monitoring and management needs.

10.1 Monitoring Storage Utilization with NetApp Operations Manager

NetApp Operations Manager monitors, manages, and generates reports on all of the NetApp storage systems in an organization. When you are using NetApp thin provisioning, NetApp recommends deploying Operations Manager and setting up e-mail and pager notifications to the appropriate administrators. With thin-provisioned storage, it is very important to monitor the free space available in the aggregates. Proper notification of the available free space means that additional storage can be made available before the aggregate becomes completely full. For more information about setting up notifications in Operations Manager, see Configuring Alarms and Managing Aggregate Capacity in the Operations Manager Administration Guide on the NetApp Support site.

For a general overview of NetApp Operations Manager, see Operations Manager on the NetApp Support site.

10.2 Monitoring and Managing of NetApp Storage on Systems Center Operations Manager (SCOM)

NetApp provides ApplianceWatch™ PRO 2.1.1 management pack, which is an enterprise-class storage monitoring application that simplifies storage management and increases tools available to SCOM administrators for NetApp storage controllers.

Following are the key features of ApplianceWatch PRO 2.1.1; with it, you can:

- Provide easy storage management for Windows and NetApp storage administrators
- Monitor all elements of IT infrastructure (single pane of glass from MMC)
- Simplify distributed NetApp storage system monitoring
- Isolate problems quickly using alerts
- Troubleshoot performance issues using performance views
- Minimize downtime, shorten time to resolution, and provide auto-recovery tools in a virtualized environment
- Rapidly provision and clone VMs using the cmdlets

NetApp ApplianceWatch PRO 2.0 includes PRO tips that enable automatic remediation of common storage issues through SCVMM. For example, administrators can choose to automatically increase the size of a volume that contains VMs if the volume's utilization rate exceeds the established threshold. In addition to \integration with SCOM and SCVMM based on ApplianceWatch, NetApp provides technology to centralize storage management and enables key storage activities to be performed by storage, server, or application administrators. This set of activities includes configuring and executing storage deduplication, data failover, VM cloning, and thin provisioning. For example, administrators can expand and shrink LUNs on the fly to address changing application workloads without affecting production environments. The ability to shrink storage in an iSCSI or FC environment is unique to the combination of Microsoft and NetApp technologies. Microsoft administrators can also set up storage policies or use existing policies that have been set up by NetApp storage administrators. Policy-based automation can apply to backups, failover, replication, and recovery to reduce errors and provide consistent execution across all major components of the solution.

10.3 Storage Growth Management

Growing Physical Disks

It is quite easy to increase the storage for the Hyper-V physical disk (SAN LUN); however, this process should be completed only when the child VMs stored on the physical disk are shut down. NetApp recommends using SnapDrive for performing this operation efficiently. Detailed instructions can be found in the SnapDrive for Windows Installation and Administration Guide. Alternatively, it can be performed using the longer, manual process outlined in the following steps.

To grow a physical disk, follow these steps:

1. Shut down the child VM hosted on the physical disk to be expanded.
2. Open FilerView (http://<controller IP address>/na_admin).
3. Select LUNs.
4. Select Manage.
5. From the list in the left pane, select the LUN that represents the physical disk to be grown.
6. In the Size field, enter the new size of the LUN and click Apply.
7. Open Server Manager on the Hyper-V server, right-click Disk Management, and select Rescan Disks.

For Windows child VM, you can use the diskpart utility to grow the file system after it has been powered on. For more information, see A Description of the Diskpart Command-Line Utility, or for a Linux child VM, you can use ext2resize to grow a file system. For more information, see GNU ext2resize.

Growing VHD Files

VHD files can be grown. However, this process requires the child VM to be powered off. Growing the VHD file is only half of the equation for increasing available storage; you must still grow the file system after the child VM boots. Note that root volumes such as C:\ in Windows and / in Linux cannot be grown dynamically or while the system is running. For these volumes, see “Growing Bootable Volumes” later in this section. For all other volumes, use native operating system tools to grow the volume.

To grow a VHD file, follow these steps:

1. Shut down the child VM.
2. In the Hyper-V manager, highlight the child VM and click Settings (on the right).
3. In the child VM settings window, select the hard drive that must be extended and click Edit. This opens the Edit Virtual Hard Disk wizard.
4. Click Expand and click Next.
5. Select the new size, click Next, and click Finish. Make sure there is enough disk space available in the physical disk where the VHD file is located (especially for fixed-size VHDs).
6. Click OK on the Settings window and power on the child VM. Remember that although you have grown the VHD, you must still grow the file system within it.
7. Follow the guidelines in “Growing a Child VM File System (NTFS or EXT3)” later in this section.

Growing Physical Disks Presented as Pass-Through Disks to a Child VM

Growing physical disks presented to the child VM (also referred to as pass-through disks) does not require the child VM to be shut down when going through the steps of growing the disk size, unlike growing VHDs. Growing pass-through disks requires only that the child VM be rebooted for the new storage to be detected by the child VM.

1. To grow a pass-through disk, follow these steps:
2. Shut down the child VM.
3. In the Hyper-V manager, highlight the child VM and click Settings (on the right).
4. In the child VM settings window, select the hard drive that must be extended and click Edit. This opens the Edit Virtual Hard Disk wizard.
5. Click Expand and click Next.
6. Select the new size, click Next, and click Finish. Make sure there is enough disk space available in the physical disk where the VHD file is located (especially for fixed-size VHDs).
7. Click OK on the Settings window and power on the child VM. Remember that although you have grown the VHD, you must still grow the file system within it.

8. Follow the guidelines in “Growing a Child VM File System (NTFS or EXT3)” later in this section.

Growing a Child VM File System (NTFS or EXT3)

When a VHD or physical disk has been increased in size, you still must grow the NTFS or EXT3 file system residing on it after booting the child VM. This process can be done by using native or freely distributed tools but cannot be done live while the child VM is running.

To grow a child OS file system, follow these steps:

1. Remotely connect to the child VM.
2. For a Windows child VM, use the diskpart utility to grow the file system after it has been powered on. For more information, see A Description of the Diskpart Command-Line Utility, or, for a Linux child VM, you can use ext2resize to grow a file system. For more information, see GNU ext2resize.

Growing Bootable Volumes

Root volumes, such as C:\ in a Windows child VM and “/” in a Linux child VM, cannot be grown on the fly or while the system is running. There is a simple way to expand these file systems that does not require the acquisition of any additional software (except for ext2resize). This process requires the VHD that has been resized to be connected to another child VM of the same operating system type by using the processes defined earlier. After the storage is connected, the hosting child VM can run the utility to extend the file system. After extending the file system, this child VM is shut down, and the storage is disconnected. Connect the VHD to the original child VM. When you boot, you can verify that the boot partition now has a new size.

Best Practice

If you are using SnapDrive to expand the storage volume, virtual machines must be correctly configured before SnapDrive for Windows is installed in the child OS. Failure to have the iSCSI interface installed prior to SnapDrive installation results in pass-through disks failing to be available to the child OS at any time.

Best Practice

The storage administrator must minimally monitor the LUN space consumed by each host, volume space, and aggregate space. The administrator must make sure alerts on volume autosize, snapshot autodelete, aggregate nearly full/full, and volume nearly full/full are received.

10.4 Adding Exclusions in the Antivirus Software

To avoid performance issues, NetApp advises running antivirus scans when the operating system is least busy. In addition to this, certain directories must be excluded from the virus scans.

Files to exclude are these:

`vmms.exe` (Hyper-V Virtual Machine Management Service)

`vmwp.exe` (VM Worker Process)

Folders and subfolders to exclude are these:

- `%systemdrive%:\ClusterStorage`
- Any folders (default and custom) containing VM configuration files (`*.vmx`, `*.bin`, `*.vsv`)
- Any folders (default and custom) containing VHDs

- Any folders containing VM Snapshot files (`*.avhd`)

Failure to exclude the files and folders from antivirus scanning activity, whether scheduled or real time, could cause not only performance issues but also corruption of the files composing the virtual machines. For more information on the topic of antivirus use with Hyper-V, refer to Microsoft KB 961804.

11 Automation

11.1 Windows PowerShell and NetApp Data ONTAP PowerShell Toolkit

Windows PowerShell is Microsoft's task automation framework, consisting of a command line shell and related scripting language built on top of and included with the .NET framework. Windows PowerShell provides full access to COM and WMI, enabling administrators to perform administrative tasks on both local and remote Windows systems. The administrative tasks are generally performed by cmdlets. A set of cmdlets can be executed as a script.

NetApp offers the Data ONTAP PowerShell toolkit for automating NetApp storage tasks. This toolkit is a collection of Windows PowerShell cmdlets for facilitating integration of Data ONTAP into the Windows environment and management tools.

Click here for more information on the Data ONTAP PowerShell Toolkit.

12 SnapManager 1.0 for Hyper-V

With the adoption of virtualization technologies, data centers have been transformed and the number of physical servers drastically reduced. Virtualization has had many positive effects, not only reducing the number of physical systems, but also reducing network, power, and administrative overhead.

In contrast to physical environments, where server resources are underutilized, fewer resources are available in virtualized environments. Although each physical server had dedicated network and CPU resources, VMs must now share those same resources, which can result in performance issues, especially while backing up the virtual environment, because many VMs are using host network and CPU resources concurrently. As a result, backups that once completed during nonbusiness hours have seen their backup window grow.

NetApp SnapManager for Hyper-V (SMHV) addresses the resource utilization issue typically found within virtual environments by leveraging the underlying NetApp Snapshot technology, thereby reducing the CPU and network load on the host platforms and drastically reducing the time required for backups to complete. SMHV can be quickly installed and configured for use in Hyper-V environments, saving valuable time during backups and allowing quick and efficient restorations, thus reducing administrative overhead.

Leveraging NetApp Data ONTAP for Hyper-V Backup, Restore, and Disaster Recovery

Backups, restores, and disaster recovery can place a huge overhead on the Hyper-V virtual infrastructure. NetApp SnapManager for Hyper-V simplifies and automates the backup process by leveraging the underlying NetApp Snapshot and SnapRestore® technologies to provide fast, space-efficient, disk-based backups and rapid, granular restore and recovery of virtual machines (VMs) and the associated datasets. The following chapters detail the best practices for deploying and using SnapManager 1.0 for Hyper-V.

12.1 Purpose and Scope

The purpose of the following chapters is to provide best practices for deploying SMHV to back up and recover Hyper-V VMs. They describe the key features and best practices to effectively manage the complete backup lifecycle for Hyper-V VMs. For detailed instructions on installation and configuration, refer to the SnapManager for Hyper-V Installation and Administration Guide.

12.2 Intended Audience

The following chapters are intended for Hyper-V administrators, storage administrators, backup administrators, and architects implementing a backup, restore, and disaster recovery solution for Hyper-V environments running on NetApp storage. Ideally, readers should have a solid understanding of the architecture, administration, and backup and recovery concepts within a Hyper-V environment and should consider reviewing the following documents:

- Data ONTAP 7.3 System Administration Guide or later
- SnapManager 1.0 for Hyper-V Installation and Administration Guide
- SnapDrive 6.3 for Windows Installation and Administration Guide or later

  For SnapManager for Hyper-V best practices on Cluster-Mode 8.1 refer to TR 4004 - Data ONTAP 8.1 operating in Cluster-Mode: Best practices for NetApp SnapManager for Hyper-V

Technical Details

SMHV provides the following capabilities:

- Allows system administrators to create hardware-assisted backup and restore of Hyper-V VMs running on NetApp storage.
- Provides integration with Microsoft Hyper-V Volume Shadow Copy Service (VSS) writer to quiesce the Hyper-V VMs before creating an application-consistent Snapshot copy of the VM.
- Allows an administrator to create application-consistent backups of Hyper-V VMs, if the customer has Microsoft Exchange, Microsoft SQL Server®, or any other VSS-aware application running on virtual hard disks (VHDs) in the VM.
- Provides mirroring of backup sets to secondary locations for disaster recovery (DR) planning.
- Supports the backup and restore of shared VMs configured using Windows failover clustering (WFC) for high availability and also on Microsoft cluster shared volumes (CSVs). SMHV makes sure that the scheduled VM backups can happen seamlessly irrespective of any VM failovers.
- Supports management of multiple remote Hyper-V parent systems from one console.
- Supports performing crash-consistent backup and restore of virtual machines in SMHV 1.1

13 SMHV Planning

Microsoft Windows 2008 Server R2 with Hyper-V role enabled offers various storage infrastructure configurations and provisioning methods. Refer to Chapters 1 through 11 of this report to determine the most appropriate choices for your environment.

13.1 Storage Considerations

SMHV supports backup and restore on CSVs. SMHV can back up only VM data stored in VHDs that reside on NetApp storage. It does not back up data on pass-through or direct-attached iSCSI disks. SMHV does not support MBR LUNs for VMs running on shared volumes or CSVs. SMHV supports LUNs created on thin-provisioned volumes and can perform backups/restores on these volumes.

14 SMHV Simplified Backup and Recovery

14.1 Prerequisites

SnapManager 1.0 for Hyper-V needs SnapDrive 6.2 for Windows (SDW 6.2) or later to be installed as a prerequisite. SnapDrive manages LUNs on a storage system, making these LUNs available as local disks on Windows Hyper-V hosts. This allows Windows hosts to interact with the LUNs just as if they belonged to a directly attached redundant array of independent disks (RAID).

Note: SDW is required on Hyper-V parent hosts, but not required on client hosts. For WFC configurations, SDW and SMHV must be installed on each node of the cluster.

Note: SMHV 1.1 supports crash-consistent backup and restore of virtual machines. This has SnapDrive 6.4.1 for Windows as a prerequisite.

14.2 Terminology

Datasets

A dataset is a grouping of virtual machines that helps you to protect data using retention, scheduling, and replication policies. You can use datasets to group VMs that have the same protection requirements. A VM could be a member of multiple datasets. This can be useful for VMs that belong to multiple groupings (for example, a VM running the SQL Server instance for a Microsoft Office SharePoint Server configuration might need to belong to both the SQL Server and the MOSS datasets).

Protection Policies

Policies allow customers to schedule/automate the backups of the datasets at a predefined time (schedule policy), allow customers to provide retention capabilities for older backups (retention policy), and allow customers to replicate the block changes to the SnapMirror destination volume after the VM backup is created (replication policy). Policy includes other capabilities that allow customers to run scripts before and after the backup.

Backup and Recovery

SMHV provides local backup and recovery capability with the option to replicate backups to a remote storage system using SnapMirror relationships.

Backups are performed on the whole dataset, which is a logical collection of VMs, with the option of updating the SnapMirror relationship as part of the backup on a per-job basis. Similarly, restores can be performed at an individual VM level.

Application-Consistent Backup/Restore

These backups are taken in coordination with the Volume Shadow copy Service (VSS) to make sure that the applications running in the VM are quiesced before creating a Snapshot copy. Such a backup guarantees the integrity of application data, and hence can be safely used to restore the VM and the applications running in the VM to a consistent state.

Crash-Consistent Backup

A backup in which the state of data is equivalent to what would be found following a catastrophic failure that abruptly shuts down the system. The data in the backup will be the same as it would be after a system failure or power outage. This type of backup is much quicker. A restore from such a backup would be equivalent to a reboot following an abrupt shutdown.

Note: Crash-consistent backup and restore are supported from SMHV 1.1 onward and will require SnapDrive for Windows 6.4.1 to be installed on the host system.

Backup Retention Policy

Retention policies can be used to specify how long you want to keep a dataset backup based on either time or number of backups. Policies can be created specifying the retention period, allowing administrators flexibility to meet varying service-level agreement (SLA) levels within their environment.

Alert Notification

Alert notifications are created on a per-scheduled-backup-job basis and are sent by e-mail to administrator-defined accounts. Alert notification can be configured to e-mail the specified account after every backup, although this is not recommended because the number of e-mails can become unmanageable. Configuring alerts to notify administrators after an error or warning within a backup offers a more useful and practical alert level.

Unprotected Resources

Unprotected resources are VMs that are not part of any dataset. These resources can be protected by adding them to a dataset.

14.3 Port Usage

Best Practice

For SMHV and SDW, make sure that the following ports are kept open:

- 808: SMHV and SDW default port
- 4094: If SDW is configured to use HTTP protocol
- 4095: If SDW is configured to use HTTPS protocol

When SMHV is installed on a cluster, the same port number must be used across all nodes.

14.4 Architecture

Figure 25 illustrates the SMHV architecture and the components that work together to provide a comprehensive and powerful backup and recovery solution for Hyper-V environments.

Figure 25) SMHV architecture.

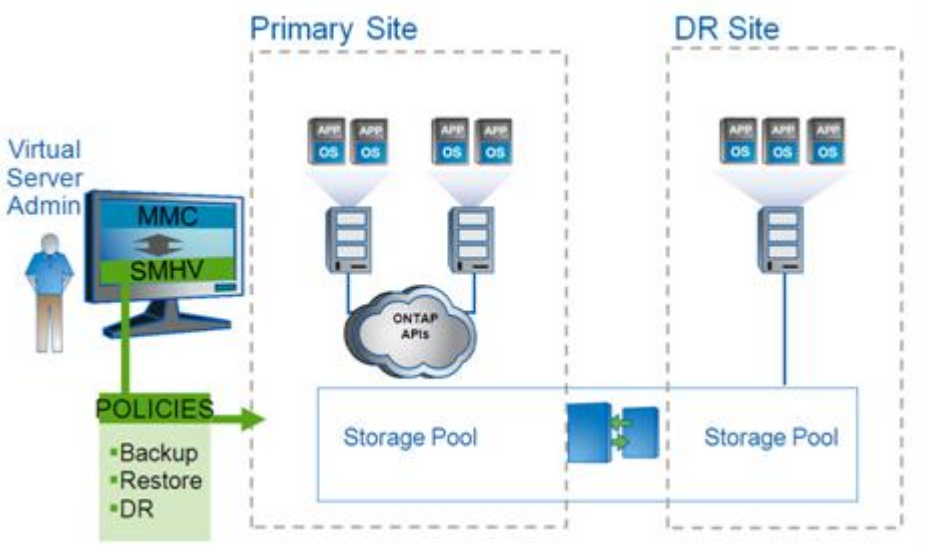

Components

License Requirements

A SnapManager for Hyper-V license is required on the Windows host system. You can choose either host-based licensing or storage system licensing.

- If you select host-based licensing, you need to provide a license key during installation. You can change the license key after installation by clicking License settings in the SnapManager for Hyper-V Welcome window.
- If you select storage system licensing, you must add the SMHV license to all storage systems.

NetApp Data ONTAP

SMHV functions only within a NetApp storage environment. SMHV requires that the primary storage where the VMs actually reside and the secondary storage used as the SnapMirror destination run the Data ONTAP storage software.

Table 15) Licensing and Data ONTAP versions.

If you use	Then use
Host-based licensing	Data ONTAP 7.3.1P1 or later
Storage system licensing	Data ONTAP 7.3.2 or later
Storage system licensing with DATA ONTAP vFiler® units	Data ONTAP 7.3.1.1P8, 7.3.2P1, or later
Storage system licensing with Cluster-Mode	Data ONTAP 8.1 or later

For the most current information, see the NetApp Interoperability Matrix Tool (IMT) at http://now.netapp.com/NOW/products/interoperability.

In addition, the following licenses are required:

- SnapRestore
- The required protocol license (FCP, iSCSI)

- SnapMirror (if required)
- SnapDrive for Windows (must be licensed on the Hyper-V host)

Configurations Supported by SMHV

SMHV must run on Windows Server 2008 R2 x64.

Platform Support

- Windows Server 2008 R2 x64 Standard, Data Center, Enterprise, Editions (Full and Core Installation)
- Hyper-V Server 2008 R2 x64

VM Support

- Windows Server 2008 R2 x64 (all editions): core and full

Windows Server 2008 x64 Standard and Enterprise Editions (full and core)

- Windows Server 2008 x64 Standard and Enterprise Editions with SP2 (full and core)
- Windows Server 2003 x64 and x86 with SP2 and later
- Windows Vista
- Windows XP
- SuSE Linux (SLES10 SP1 and SP2) x86 and x64
- RHEL 5.3, RHEL 5.4, and RHEL 5.5 (Microsoft Hyper-V Integration component version 2.1 must be installed)

For the most current information, see the NetApp IMT at http://now.netapp.com/NOW/products/interoperability.

SMHV SnapInfo Settings

The SMHV SnapInfo folder stores backup metadata. This can be set up by specifying the SnapInfo settings in the Hosts Management wizard. The metadata information is critical to recovering VMs should a failure occur. SnapInfo settings should be configured for the host or cluster added to SMHV so that VMs within that host can be added to a dataset.

Note: The SnapInfo path must reside on a Data ONTAP LUN. For managing dedicated VMs, the SnapInfo location must be a dedicated Data ONTAP LUN. For managing shared VMs, the SnapInfo location must be to a shared Data ONTAP LUN.

The SnapInfo path must not reside on a CSV.

Note: If SnapInfo settings are changed, you must manually move all files from the original SnapInfo location to the new location. SnapManager for Hyper-V does not move them automatically.

Best Practice
NetApp recommends having the SnapInfo LUN on a volume of its own.

SMHV Report Settings

Report settings should be configured for a host or cluster added to SMHV so that VMs within that host can be added to a dataset.

Best Practice
The report path must not reside on a CSV.

SMHV Event Notifications

Event notifications settings can be configured to send e-mail and AutoSupport messages in case an event occurs.

15 SMHV Process Flow

15.1 Adding a Hyper-V Parent Host or Host Cluster

If you add a single host, SMHV manages the dedicated VMs on that host. If you add a host cluster, SMHV manages the shared VMs on the host cluster. If you plan to add a host cluster, SMHV must be installed on each cluster node.

If the backup repository settings, report directory settings, and notification settings are not configured for SMHV, you can configure them after you add the host, using the configuration wizard. You must configure the backup repository and report directory settings to add and manage VMs using SMHV. Notification settings are optional.

Note: Dedicated and shared VMs that belong to the same host cluster should not exist in the same dataset. Adding these types of resources to a single dataset can cause the dataset backup to fail.

Although you should manage a host from only one management console, if you need to do so from multiple consoles, you can import and export host and dataset configuration information from one remote management console to another to be sure of data consistency. You can also use the Import and Export wizard to change host and dataset configuration settings to a previously exported setting. If you perform this operation in a clustered environment, you must import the settings on all nodes in the cluster so that all host and dataset configurations are the same. You should not import or export configuration information to the directory where SMHV is installed. If you uninstall SMHV, this file will be lost.

The Backup Process and Implications

SMHV leverages NetApp Snapshot technology to create fast and space-efficient backups of SMHV datasets and their associated VMs. These backups offer point-in-time images, or copies, of the VMs and are stored locally on the same storage platform on which the VMs physically reside.

In addition to the Snapshot copy stored locally, SMHV also provides an option to update an existing SnapMirror relationship upon the completion of a backup. This can be selected on a per-backup-job basis as required by the administrator. The unit of backup in SMHV is a dataset, which can contain one or more VMs running across multiple Hyper-V hosts. SMHV supports restoring an individual VM; it does not support restoring an entire dataset.

Using SMHV, on-demand or scheduled backups of VMs are possible. SMHV supports backup of dedicated or clustered VMs. It also supports backups of shared VMs running on CSVs.

Figure 26) Hyper-V infrastructure and associated storage during an application-consistent SMHV backup.

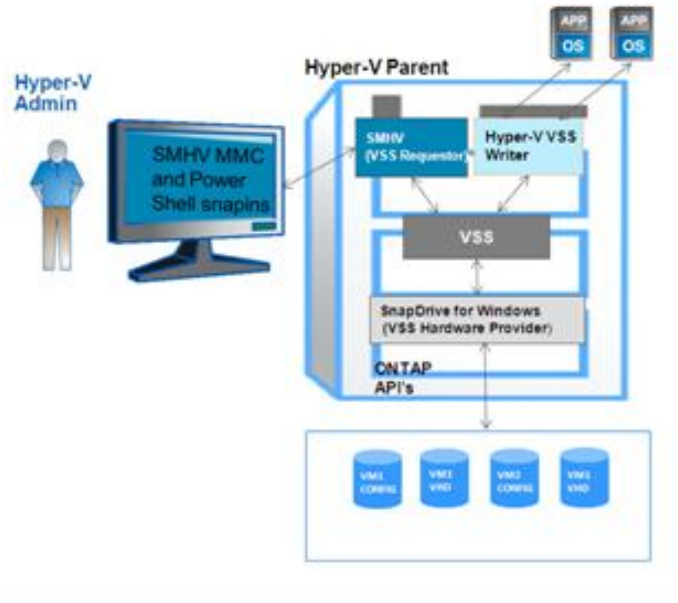

Figure 26 represents a high-level overview of the typical SMHV architecture on the primary site storage and will be used in detailing the backup process flow for application-consistent backups.

1. The SMHV service is a VSS requestor that initiates a VSS backup of VMs within a dataset in coordination with the Microsoft Hyper-V VSS writer.
2. The Hyper-V VSS writer works together with the integration services within the VM to create application-consistent “software” Snapshot copies of all VHD volumes attached to each VM.
3. SMHV then implements a VSS requestor component to coordinate the backup process and create a consistent Snapshot copy in Data ONTAP using a VSS hardware provider for Data ONTAP LUNs.
4. VSS framework requests the hardware provider to mount the LUNs from the Snapshot copy.
5. Hyper-V writer recovers data on the LUNs and brings it to the state of the software Snapshot copy that was created in step 2.
6. The VSS provider creates a second Snapshot copy of the LUNs and then dismounts them from the Snapshot copy.
7. Upon completion of the local backup, SMHV updates an existing SnapMirror relationship on the volume if the SnapMirror option was selected. SnapMirror will be discussed in further detail in a later section of this document.

SMHV enables you to create application-consistent backups of a VM, if you have Microsoft Exchange, Microsoft SQL, or any other VSS-aware application running on VHDs in the VM. SMHV coordinates with the application VSS writers inside the VM to make sure that application data is consistent when the backup occurs.

Note: For a backup to succeed, all files of the VM (VHDs, VM configuration files, and VM Snapshot files) should reside on LUNs managed by Data ONTAP.

Note: Only one backup operation can occur on a host at any given time. If the same VMs belong to different datasets, you should not schedule a backup of the datasets at the same time. If this occurs, one of the backup operations will fail.

Note: SMHV backup fails for VMs that have a VHD created by copying the contents of a physical disk on the same host. The Create New VHD wizard of Hyper-V manager gives this option. As part of copying the physical disk contents, it also copies the disk signature, and this causes the disk

signature conflict during the backup. More information is available here: http://support.microsoft.com/kb/975695.

Do not create a VHD using the option “copy the contents of the specified physical disk” in the “configure disk” page in the new VHD creation wizard in Microsoft Hyper-V manager.

SnapManager for Hyper-V does not support the backup and restore of virtual machines running on SAN boot LUNs. This is a limitation of SDW.

Workflow for Crash-consistent backups:

- User chooses crash-consistent backup option in the backup dataset wizard.
- SMHV API calls VSS to collect the VM metadata. The LUNs on which the VMs are hosted are identified.
- The SnapDrive API is called to take a Snapshot copy of these LUNs. Only one Snapshot copy will be taken for each LUN irrespective of the number of VMs running on it.
- Backup will be registered with backup type as 'Crash-consistent.'
- Upon completion of the local backup, SMHV updates an existing SnapMirror relationship on the volume if the SnapMirror option was selected.

Note: While performing a crash-consistent backup or restore, SMHV 1.1 does not leverage VSS. VSS is used only to get VM-related metadata from the Hyper-V writer. The default backup type will be application-consistent backup.

Best Practice
When creating a dataset, you should select all VMs that reside on a particular Data ONTAP LUN. This enables you to get all backups in one Snapshot copy and to reduce the space consumption on the storage system. It is preferable to add VMs running on the same CSV in the same dataset. If you add VMs on the same CSV in different datasets, make sure that the backup schedules of these datasets do not overlap.

Best Practice
If you change a VM Snapshot copy location to a different Data ONTAP LUN after creating the VM, you should create at least one VM Snapshot copy using Hyper-V manager before creating a backup using SMHV. If this is not done, the backup could fail.

15.2 Scheduled Backups and Retention Policies

SMHV allows administrators to schedule a dataset backup at a particular time. SMHV uses the Windows Tasks Scheduler for creating or modifying scheduling policies. The 255 NetApp Snapshot copies-per-volume limit must be taken into consideration when scheduling backups and configuring the associated retention policies. The number of Snapshot copies per volume can be managed with the proper scheduling and retention policies on a per-scheduled-backup basis while still meeting SLAs on the VMs.

Backup Scheduling

Using scheduling policies, administrators can schedule backup jobs at particular times, allowing them to automate the process. Multiple policies can be scheduled per dataset that apply to all hosts that are dataset members.

Best Practice

The backup frequency, as well as the number of different backups performed against a dataset—for example, one backup running against dataset ds_1 weekly and another monthly—must be taken into account when specifying the retention policy so as not to exceed the maximum number of Snapshot copies per volume. Should the number of Snapshot copies exceed 255 on any given volume, future backups against that volume will fail.

Retention Policies

The following list describes the retention tags available in SMHV:

- **Hourly.** Hourly intervals
- **Daily.** A specified time within a 24-hour period
- **Weekly.** A specified day and time within a seven-day period
- **Monthly.** A specified day and time within a calendar month
- **Unlimited.** Never-deleted backups

After choosing a retention type, you can choose to delete either backups that are older than a specified period of time or backups that exceed a maximum total.

NetApp recommends using the policies not only to meet specific SLAs, but also to maintain a supported number of NetApp Snapshot copies on the underlying volumes. For SMHV, one backup creates two Snapshot copies on the storage systems for data consistency (refer to KB ID 2010607). For example, setting a retention policy of 30 backups on an hourly basis limits the maximum number of Snapshot copies associated with the backup to 60. However, if the retention policy had been configured as 30 days, the Snapshot limit per volume would be reached in 5 days, and backups would begin to fail from that point on.

Best Practice

Choose a backup retention level based on your backup creation and verification schedule. If a Snapshot copy deletion occurs, make sure that a minimum of one verified backup remains on the volume. Otherwise, you run a higher risk of not having a usable backup from which to restore in case of a disaster.

Note: The option “unlimited” should be used with caution. When this option is selected, backups and the associated NetApp Snapshot copies are maintained until they are manually deleted by the administrator. These Snapshot copies are included in the maximum number supported on a volume.

Of further note, the NetApp Snapshot copies associated with on-demand backups must also be considered when determining the number of Snapshot copies maintained against a volume.

After creating a dataset backup, SMHV creates a Snapshot copy of the SnapInfo LUN. SnapInfo Snapshot copies are not deleted if the backup is deleted. SnapInfo Snapshot copies have a different retention policy. By default, SMHV retains 30 SnapInfo LUN Snapshot copies and deletes the older ones when the SnapInfo Snapshot count exceeds 30. You can configure the number of SnapInfo Snapshot copies you want to retain for each Hyper-V host using the following registry key:

- For standalone Hyper-V hosts:
  Registry key: `HKLM\SOFTWARE\NetApp\SnapManager` for Hyper-V\Server DWORD value:
  `snapinfo_snaps_count` (number of SnapInfo Snapshot copies to be retained)
- For clustered Hyper-V hosts (to be configured on each node in the cluster):
  Registry key: `HKLM\Cluster\SOFTWARE\NetApp\SnapManager` for Hyper-V\Server DWORD

value:
snapinfo_snaps_count (number of SnapInfo Snapshot copies to be retained)

15.3 Handling Saved-State Backup of VMS

The default behavior of SMHV is to fail a backup if one or more VMs cannot be backed up online. If a VM is in the saved state or shut down, an online backup cannot be performed. In some cases, VMs are in the saved state or shut down for maintenance, but backups still need to proceed, even if an online backup is not possible. To do this, the VMs that are in the saved state or shut down can be moved to a different dataset with a policy that allows saved-state backups.

Note: You can also select the Allow saved-state VM backup checkbox to allow SMHV to back up the VM using the saved state. If you check this option, SMHV will not fail the backup when the Hyper-V VSS writer backs up the VM using the saved state or performs an offline backup of the VM. Doing a saved state or offline backup can cause downtime. For more information on online or offline VM backups, see the Hyper-V Planning for the Backup information in the Technet library:http://technet.microsoft.com/en-us/library/cc753637(WS.10).aspx.

Best Practice
For mission-critical VMs, NetApp recommends disabling the “Allow Saved state VM backup” option.

Note: ‘Allow saved state policy’ option is not applicable for crash-consistent backups. This is because the VM is being backed up irrespective of the state.

15.4 Backup Scripts

Using SMHV, you can run optional backup scripts either before or after the backup takes place. These scripts will run on all dataset member hosts unless you indicate a specific server. The following environment variables can be used as arguments for application-consistent backup postscripts:

- `$VMSnapshot`

  Specifies the first VM Snapshot copy name that is created on a storage system as a result of the backup. The second name uses the first name plus the `appendix _backup`.

- `$SnapInfoName`

  Specifies the time stamp used in the SnapInfo directory name.

- `$Snapinfosnapshot`

  Specifies the SnapInfo Snapshot copy name created on the storage system. SMHV makes a Snapshot copy of the SnapInfo LUN at the end of the dataset backup.

During the post–script execution phase, SMHV replaces the $VMSnapshot variable with the Snapshot name, `$SnapInfoName` with the time stamp of the backup, and `$SnapInfoSnapshot` with the SnapInfo Snapshot name.

Note: The $SnapInfoSnapshot variable is supported for dedicated virtual machines only.

15.5 Quick/Live Migration Implications

Best Practice
SMHV cannot back up a VM that is actively undergoing migration. When a backup runs against a dataset that has VMs actively being migrated, an error is generated, and those particular VMs are not backed up.

15.6 Restore Process

SMHV can restore a VM from a backup. SMHV can also restore a VM that is part of a cluster. To restore the VM, SMHV uses the file-level restore feature in SDW. You can spread the associated files of a VM, including the configuration file, Snapshot copies, and any VHDs, across multiple Data ONTAP LUNs. A LUN can contain files belonging to multiple VMs.

If a LUN contains only files associated with the VM you want to restore, SMHV restores the LUN using LUN clone split restore (LCSR). If a LUN contains files not associated with the VM you want to restore, SMHV restores the VM using the file copy restore operation.

With these differences in restore types aside, the process flow used by SMHV during a restore is as follows:

1. SMHV restores a VM in coordination with Hyper-V VSS writer. Hyper-V VSS writer powers off the VM and deletes it before restore.
2. Files are restored as described in the preceding paragraphs based on restore type.
3. SMHV notifies the VSS writer that the files of the VM are restored properly. Hyper-V VSS writer registers the VM, and the VM gets added back in the Hyper-V manager.
4. SMHV starts the VM after restore and executes a postscript if specified in the restore wizard.

Note: During the restore, the following warning messages might be displayed:

- VM to be restored is not [currently running] on the host.
- VM to be restored is currently running on the host, and:
  - It has more VHDs associated with it than at the time of backup.
  - It has fewer VHDs associated with it than at the time of backup.
- The Snapshot location of the VM has changed.
- The names of VHD files or their file system paths or NetApp storage system LUN paths have changed.

  In all of these warning scenarios, the VM can be restored, but you must acknowledge that you are sure you want to go ahead with the restore.

Note: If the VM no longer exists, you can still restore it if the LUNs on which the VM was created still exist. The LUNs must have the same drive letters and Windows volume GUIDs as at the time of backup.

If the VM no longer exists, you can still restore it by selecting a backup to which it belonged.

If the VM was removed from all datasets before it was deleted, you can still restore it by selecting unprotected resources and selecting a backup to which it belonged.

Best Practice

If the number of VHDs attached to a VM at the time of backup and restore is not same, the restored VM might have additional/fewer VHDs. If that is the case, NetApp recommends that the cluster configuration of the VM and its dependencies be manually updated.

Note: SMHV does not back up the cluster configuration of the VM, so it does not restore the cluster configuration. If the VM and the cluster configuration are lost, you can restore the VM from SMHV, but you must manually make it highly available. For more information, see "Failover Clustering on Windows Server 2008 R2" on the Microsoft Web site.

Note: In case of crash-consistent backups, the VM is restored without involving the VSS. It performs a file level restore of the VM using SnapDrive for Windows.

Note: Restoring a deleted VM is not supported for Crash-consistent backups. Also, RestoreToAlternateHost switch in Restore-Backup cmdlet cannot be used when the backup being restored is a crash-consistent backup.

15.7 Mounting a Backup

Backups can be mounted using SnapDrive for Windows. The mounted backup is a clone of the protected VM. Once mounted, the backup is displayed within the explorer of Hyper-V host and can be browsed.

1. Select the LUN, and within Snapshot copies select the backup to mount.

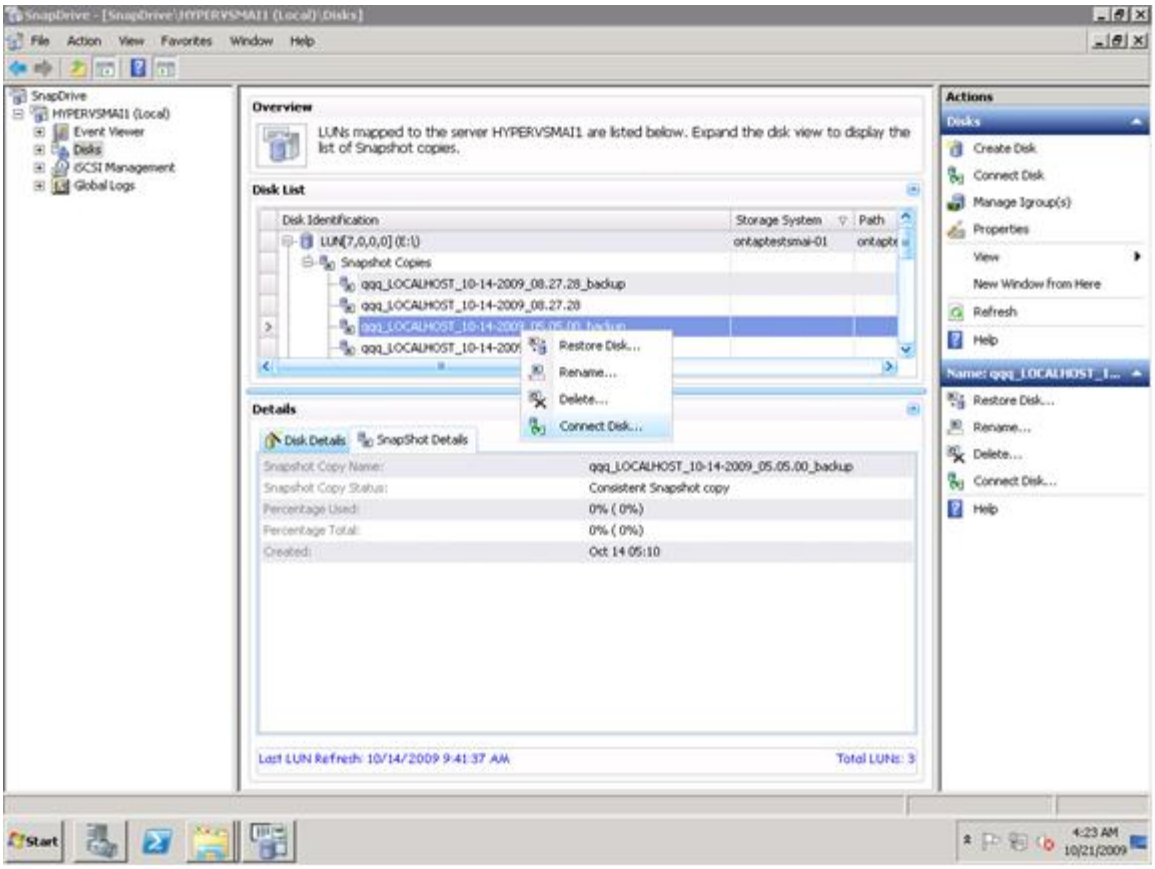

2. Right-click the Snapshot copy (the one with _backup suffix) and select the connect disk option.
3. Click Next.
4. If the LUN is a dedicated disk, go to the next step; otherwise, if the LUN is a Windows cluster resource, perform the following steps in the Specify Microsoft Cluster Services Group panel. In the Specify Microsoft Cluster Services Group panel, perform one of the following actions and then click Next.
   a. Select a cluster group from the Group Name drop-down list.
   b. Select Create a new cluster group to create a new cluster group.

Note: When selecting a cluster group for your LUNs, choose the cluster group your application will use.

Note: If you are creating a volume mount point, the cluster group is already selected. This is because the cluster group owns your root volume physical disk cluster resources. NetApp recommends that you create new shared LUNs outside of the cluster group.

   c. Select Add to cluster shared volumes.

5. In the Select LUN Properties panel, perform the following actions: Either select a drive from the list of available drive letters, or enter a mount point for the LUN you are connecting. When you create a volume mount point, enter the drive path that the mounted drive will use; for example, G:\mount_drive1\.
6. In the Select Initiators panel, choose an initiator for the LUN.
7. In the Select Initiator Group Management panel, specify whether you will use automatic or manual igroup management.
8. In the Completing the Connect Disk Wizard panel, perform the following actions:
   a. Verify all the settings.
   b. If you need to change any settings, click Back to go back to the previous wizard panels.
   c. Click Finish.
9. Browse the backup by selecting the drive letter on the explorer of Hyper-V host.

Single-File Restore Capability

In addition to backup verification, mounting a backup provides a way to restore a single file from within a VM on a case-by-case basis. This is performed by attaching a VHD from within the mounted backup as an existing hard drive to a VM within Hyper-V manager. Once a backup has been mounted, the user can use the “Hot Disk Add” functionality in Windows 2008 R2 to attach a disk (backed by the VHD) to the VM at run time without shutting down the VM. Using this functionality, the user can attach new disks to the VM.

This is a three-step process, as the following procedure shows:

1. The user must first mount the VHD from the backup mounted location [<drive>:\ Name.vhd] to the parent host using the Attach VHD option from Disk Management SnapIn. This mounts the VHD as a new disk in the Hyper-V parent.
2. Offline the disk just mounted in the preceding step using the Disk Manager Snapin. Select the disk and choose the offline menu item. This offlines the disk mounted from VHD.
3. Attach the offlined disk to the virtual machine by selecting the Physical hard disk radio button and choose the disk that just offlined from the VM settings property page, as shown in the following screen capture.

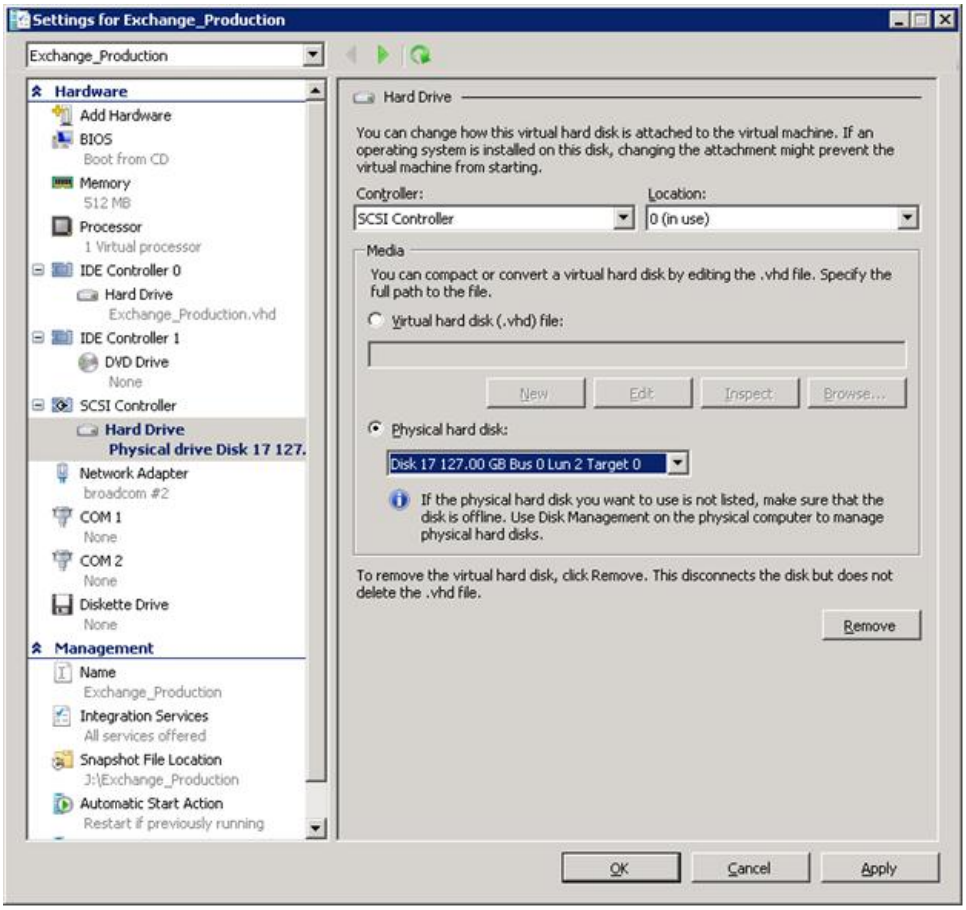

This presents a new drive inside the VM (backed by VHD in the parent). The user can then log in to the VM and choose the newly mounted drive and see the contents of the disk backed by the VHD attached.

Once the verification is done, detach the disk from the virtual machine using the VM settings page and choose the Remove button. Use SnapDrive for Windows to unmount the disk backed by the Snapshot copy, using the SnapDrive Disconnect disk MMC action/menu item. Customers can also use SDCLI Snap Unmount command to unmount the disk mounted from Snapshot technology.

Note: Leaving a backup in a mounted state places Snapshot copies in a busy condition, preventing the deletion of both the mounted backup and any preceding Snapshot copies. Backup should be unmounted when not in use.

16 SMHV High Availability

The availability of the shared storage infrastructure is more critical than the actual availability of the individual physical servers hosting the VMs on a Hyper-V server itself as they support features such as live/quick migration, which makes sure of the high availability at the hypervisor layer. With the NetApp software solution, most of the availability requirements of a virtual infrastructure can be addressed.

Note that the SMHV, as a host-end application, offers services provided that the storage is continuously available. Following is a detailed description of the available tools that facilitate storage availability.

16.1 Multipath HA with Active-Active NetApp Controllers

The NetApp active-active controllers offer easy, automatic, and transparent failover capabilities to deliver a high-availability (HA) solution. Configuring multipath HA with NetApp active-active controllers enhances the overall storage infrastructure availability and promotes higher performance consistency. It offers protection against storage failure events such as FC adapter or port failure, controller-to-shelf cable failure, shelf module failure, dual intershelf cable failure, and secondary path failure. This equips environments running business-critical applications such as the Microsoft Hyper-V virtual infrastructure to provide uninterrupted services.

Best Practices
Use active-active storage controller configuration to eliminate any single points of failure (SPOFs).
Use multipath HA with active-active storage configuration to get a better storage availability and higher performance.
More details on high-availability system configuration can be obtained from NetApp TR-3450: Active-Active Controller Overview and Best Practices Guidelines.

16.2 Data ONTAP DSM for Windows MPIO

Microsoft MPIO is a protocol-independent feature that supports multiple data paths to a storage device with iSCSI, Fibre Channel, or SAS. Providing multiple paths that can handle failover increases the availability from a host to the storage system. Windows 2008 R2 x 64 servers include support for Microsoft MPIO.

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

Best Practices
For a highly available connection to the storage system, NetApp requires installing the supported version of multipathing software such as the Data ONTAP DSM for Windows MPIO.
For Windows Server 2008 R2 servers, NetApp recommends Data ONTAP DSM 3.2R1 or higher.
For the currently supported multipathing software versions and related requirements, see the NetApp Interoperability Matrix.

17 SMHV Disaster Recovery

The disaster recovery functionality can be used after updating the SnapManager for Hyper-V software with a patch. The patch is available at http://now.netapp.com/NOW/download/software/snapmanager_hyperv_win/1.0P1/.

After upgrading to SMHV 1.0 P1, the user can perform failover and failback of Hyper-V VMs using Windows PowerShell cmdlets in the SMHV PowerShell option. The Windows PowerShell cmdlet `restore-backup powershell cmdlet` must be used along with the switch `–RestoreToAlternateHost` and the server name to use this feature.

For example:

```
PS C:\Windows\system32> restore-backup -server cluster_1 -RestoreToAlternateHost -
disableverifysnapshot -backup DR_Dataset_Secondary_01-22-2010_18.21.33 -resourcename smhv-demo-
csv -verbose
```

17.1 New Cmdlet: Get-VMsFromBackup

This cmdlet is used to retrieve the VMs from backup metadata. In a DR scenario, the administrator has access to the backup metadata from primary. The administrator needs to know which VMs are present in the backup to be able to restore them on secondary. This new cmdlet provides a list of VMs present in the backup.

The `-server` switch of this cmdlet is used to specify the hostname or cluster name on the secondary site. SMHV looks for the backups in SnapInfo for this input host/cluster and finds out VMs present in these backups.

For example:

```
PS C:\Windows\system32> get-vmsfrombackup -server cluster_windows2008_r2
Name Id
SMHV-demo-CSV F10F1011-901A-4789-ADE4-A1F34323E2D7
```

17.2 Basic DR Scenario

Components:

- Site A (primary) containing storage systems and standalone Hyper-V host system or Hyper-V host cluster. VMs running on these hosts are residing on NetApp storage.
- Site B (secondary) containing storage systems and Hyper-V host or cluster (same as that of primary).
- SnapDrive for Windows and SnapManager for Hyper-V are installed on both site A and site B.
- SnapMirror relationship is initialized from site A to site B.
- Hyper-V host or cluster on site A is added to SMHV, and the VMs are being backed up using SMHV. The policy to update SnapMirror after backup is checked. So, after each backup, the secondary site is updated with new Snapshot copies of VMs and SnapInfo.

Steps to Fail Over VMs to Secondary

1. Connect to all the LUNs from secondary storage system volumes. If secondary is a cluster, go to the node where cluster group is online and connect to all the LUNs from that node in the cluster. The LUN type and mount point must be the same as that of the primary. SDW breaks the SnapMirror relationship and also does SnapRestore. If the volume contains only one LUN, SDW performs a volume-based SnapRestore (VBSR), and the SnapMirror relationship is then in uninitialized state. If the volume contains multiple LUNs, SDW performs a single–storage system SnapRestore (SFSR), and the SnapMirror relationship is broken off.
2. Restore the SnapInfo LUN from its last Snapshot copy created by SMHV.
3. Add the secondary host or cluster in SMHV and configure it with the SnapInfo path.
4. Use the Get-VMsFromBackup cmdlet to get list of VMs present in backup metadata.
5. Use the Get-Backup cmdlet to get the backups for each VM.
6. Use Restore-backup cmdlet with VM GUID (from step 5) and backup (from step Use -RestoreToAlternateHost switch and specify the secondary host or cluster name as -server parameter. If secondary is a cluster, make sure that the LUNs on which VMs reside are online on the cluster node that owns the cluster group.
7. If secondary is cluster, make VMs highly available using failover cluster UI/Windows PowerShell cmdlets.

Note: If the secondary site is an active site with its own virtual machine LUNs and SnapInfo LUN, then in order to restore the VMs present in the primary site to the secondary site:

1. Connect the primary SnapInfo LUN to the secondary host by breaking the mirrored volume.
2. Snap restore from the last SMHV Snapinfo Snapshot copy.

3. Copy the contents to the already existing SnapInfo copy to the secondary.

In this manner, the VMs in the primary are reflected in the SMHV console of the secondary site and can be managed appropriately.

Steps to Fail Back VMs to the Primary

1. Get the data from secondary back on primary storage system.

   If primary site is completely destroyed, new storage has to be provisioned. If that is done, the user must initialize the SnapMirror relationship from secondary to primary (this is a new relationship) to get the data back. After the relationship is initialized and the data is back on primary, this relationship can be released. If primary site was temporarily down, the user must get only those changes to primary that happened on secondary while primary was gone. To do this, resync the existing SnapMirror relationship in reverse direction (resync from secondary to primary).

2. When the data on the secondary is synchronized with primary, go to SnapDrive UI on secondary and initiate a SnapMirror update for each of the LUNs on the secondary. If this is not done, SDW uses the SMHV backup Snapshot copy to restore the LUNs on primary during connect in step 3. The LUN in the backup Snapshot copy is actually a LUN clone, so this must be avoided by forcing one more SnapMirror update.

   Taking SMHV backup (with the `Update SnapMirror` option checked) from the secondary has the same effect as manually doing the SnapMirror update from SDW GUI. Most users will probably take the SMHV backup in lieu of manually doing mirror update because it can be scripted, whereas the mirror update is a tedious job and prone to user error (such as forgetting to update a LUN).

3. Connect to all LUNs on primary (same type, same mount points). If primary is a cluster, go to the node where the cluster group is online and connect to all the LUNs from that node in the cluster. If a resync in reverse direction has been done, there will be a new broken (or uninitialized) SnapMirror relationship from secondary to primary. This can be released.
4. Restore the SnapInfo LUN from its last Snapshot copy created by SMHV.
5. Add the primary host or cluster in SMHV MMC and configure it with the SnapInfo path.
6. Use the Get-VMsFromBackup cmdlet to get list of VMs present in backup metadata.
7. Use the Get-Backup cmdlet to get the backups for each VM.
8. Use Restore-backup cmdlet with VM GUID (from step 6) and backup (from step 7). Use -RestoreToAlternateHost switch and specify the primary host or cluster name as -server parameter. If primary is a cluster, make sure the LUNs (cluster resources) on which the VM resides are online on the node that owns the cluster group.
9. If primary is the cluster, make VMs highly available using failover cluster UI/Windows PowerShell cmdlets.

   After the VMs are backed up on the primary, the user must get back to the original configuration with a SnapMirror relationship established from primary to secondary. To do this, perform the following steps on secondary:

10. If secondary is a standalone host, shut down and delete the VMs running on secondary. Disconnect the SnapInfo disk and the disks containing VMs using SnapDrive. If secondary is a cluster, offline the virtual machine resource and virtual machine configuration resource for all the VMs. Delete these resources from the cluster. Delete all the VMs from Hyper-V manager. Disconnect all the disks using SnapDrive.
11. Resync the SnapMirror relationship from primary to secondary.

18 SMHV Application Consistency

Microsoft's Volume Shadow Copy Service, or VSS, was written specifically to enable third-party backup and recovery solutions to provide application-consistent backup and recovery for mission-critical

applications supported by Microsoft. When VSS is properly configured within the Hyper-V environment, a Snapshot copy initiated by SMHV begins the VSS process.

VSS is designed to produce fast, consistent Snapshot copy–based online backups by coordinating backup and restore operations among business applications, file system services, backup applications, fast recovery solutions, and storage hardware. VSS coordinates Snapshot copy–based backup and restore and includes these additional components:

- **VSS requestor.** The VSS requestor is a backup application, such as the SMHV application or NTBackup. It initiates VSS backup and restore operations. The requestor also specifies Snapshot copy attributes for the backups it initiates.
- **VSS writer**. The VSS writer owns and manages the data to be captured in the Snapshot copy. Microsoft Hyper-V is an example of a VSS writer.
- **VSS provider.** The VSS provider is responsible for the creation and management of the Snapshot copy. A provider can be either a hardware provider or a software provider: A hardware provider integrates storage array–specific Snapshot copy and cloning functionality into the VSS framework. The Data ONTAP VSS hardware provider integrates the SnapDrive service and storage systems running Data ONTAP into the VSS framework. A software provider implements Snapshot copy or cloning functionality in software that is running on the Windows system.

The coordinated backup process includes freezing the data application I/O, flushing the file system cached I/O to disk, and creating a point-in-time Snapshot copy of the data state. After the Snapshot copy is created, file system and application I/O is resumed. The VSS restore process involves placing the data application into the restore state, passing backup metadata back to the application whose data is being restored, restoring the actual data, and signaling the data application to proceed with recovering the data that was restored.

SMHV provides integration with Microsoft Hyper-V VSS writer to quiesce a VM before creating an application-consistent Snapshot copy of the VM. SMHV is a VSS requestor and coordinates the backup operation to create a consistent Snapshot copy, using VSS hardware provider for Data ONTAP. SMHV allows you to create application-consistent backups of a VM if you have Microsoft Exchange, Microsoft SQL Server, or any other VSS-aware application running on VHDs in the VM. The applications that exist in the VM restore to the same state as at the time of the backup. SMHV restores the VM to its original location. If applications are running on pass-through or direct-attached iSCSI LUNs, these LUNs are ignored by the VSS framework in the VM, and SMHV does not create a backup of these LUNs in the VM. To enable backup of application data on direct-attached iSCSI LUNs or pass-through LUNs in the VM, you would need to configure application backup products in the VM (for example, SnapManager for Exchange, SnapManager for SQL Server, and so on).

Note: The Data ONTAP VSS hardware provider is installed automatically as part of the SnapDrive software installation.

Note: To make sure the Data ONTAP VSS hardware provider works properly, do not use the VSS software provider on Data ONTAP LUNs. If you use the VSS software provider to create Snapshot copies on a Data ONTAP LUN, you will be unable to delete that LUN using the VSS hardware provider.

Note: VSS requires that the provider initiate a Snapshot copy within 10 seconds. If this time limit is exceeded, the Data ONTAP VSS hardware provider logs event ID 4364. This limit could be exceeded due to a transient problem. If this event is logged for a failed backup, retry the backup.

Note: SMHV coordinates with Hyper-V VSS writer to create application-consistent backup of VMs. Hyper-V writer communicates with integration services (Hyper-V Volume Shadow Copy requestor service) installed in the VM to quiesce the applications running in the VM before creating a backup. Data ONTAP VSS hardware provider installed on the Hyper-V host as part of SnapDrive is used to create Snapshot copies on storage system.

Note: For details on VM backup, refer to the following TechNet link:

Note: http://technet.microsoft.com/en-us/library/dd252619(WS.10).aspx.

19 Crash-Consistent Backup and Restore

Backups taken using SMHV 1.1 can be either application-consistent or crash-consistent. Application-consistent backups are taken in coordination with Volume Shadow copy Service (VSS) to make sure that the applications running in the VM are quiesced before taking the Snapshot copy. Such a backup guarantees the integrity of application data, and hence can be safely used to restore the VM and the applications running in the VM to a consistent state.

Though application-consistent backups are perfect solution for data protection and recovery of Hyper-V VMs, they also have a few drawbacks:

1. Application-consistent backups are slower due to VSS involvement with the parent and guest OS. As application writers in the VM and Hyper-V writer in the parent OS are involved, the backup process is more complex and hence more error prone. Failure in any of the components will fail the backup.
2. Hyper-V writer uses the auto-recovery process to make the VMs consistent. Auto-recovery results in the creation of two Snapshot copies on storage system. Therefore, each Hyper-V backup requires two Snapshot copies to be created per storage system volume.
3. If multiple VMs are running on different nodes in a cluster, but on the same CSV, SMHV still needs to create one backup per node as required by VSS. As a result, SMHV creates multiple Snapshot copies on the same CSV for different VMs.

Given these drawbacks, it will be desirable to have some way of taking "quick" Hyper-V VM backups. Crash-consistent backup is designed to provide this ability of taking quick backups.

A crash-consistent backup of a VM will not use VSS to quiesce data, nor will it result in auto-recovery. This backup will simply take a Snapshot copy on the NetApp storage system for all the LUNs used by the VMs involved in the dataset. The data in the backup will be the same as it would be after a system failure or power outage. All the SMHV functions such as scheduling, restore, script execution, SnapMirror updates, backup retention, and so on will be supported for crash-consistent backups as well.

Figure 27) Backup dataset wizard showing backup types: application-consistent and crash-consistent.

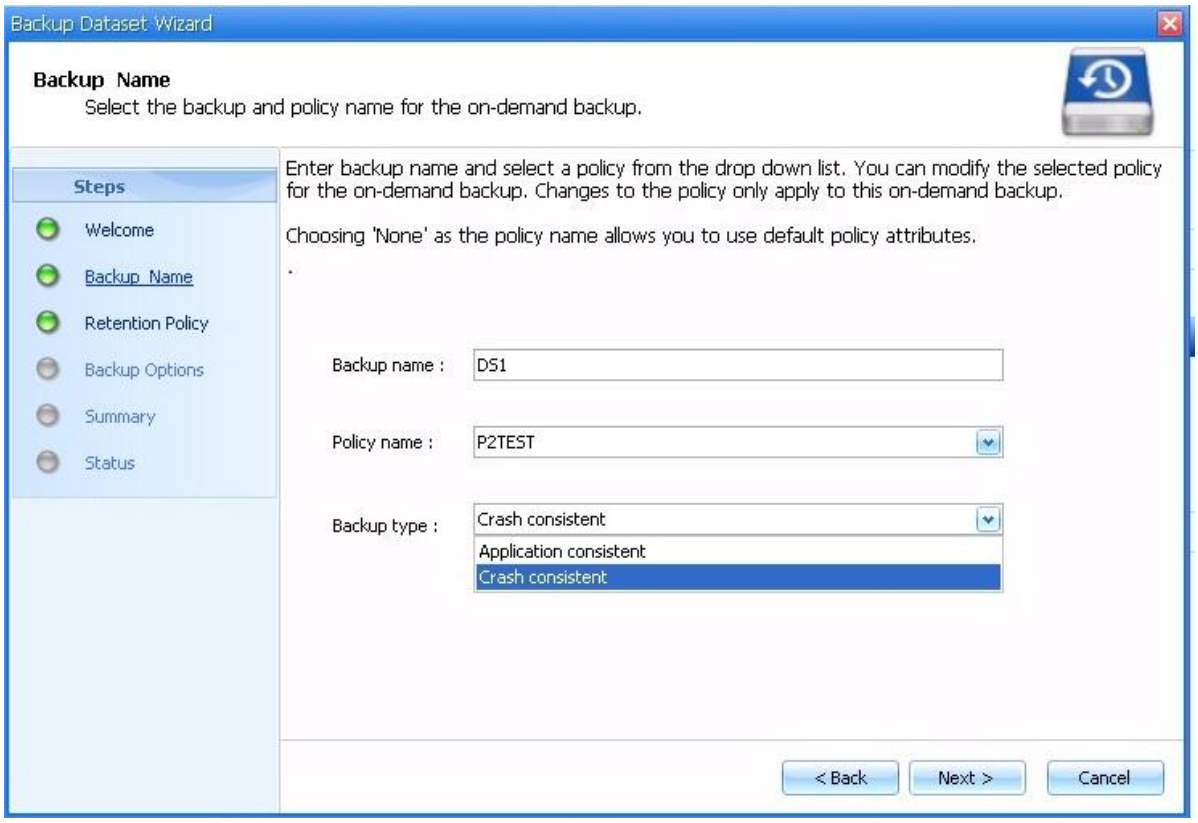

Note: Saved state backup policy is not applicable for crash-consistent backup and restore. This is because crash-consistent backups do not involve the Hyper-V VSS writer.

Note: SMHV supports parallel execution crash-consistent and application-consistent backups. It also supports parallel crash-consistent backup execution. However, users might observer some issues while such operations are executed. This is due to a timeout error in the underlying SnapDrive for Windows.

Best Practice

The crash-consistent backup feature is not a replacement for application-consistent backups. It enables user to have frequent recovery points. Therefore, user can have frequent crash-consistent backups and fewer application-consistent backups.

Best Practice

Crash-consistent backup can be used to take the latest backup of all the data immediately before performing an application-consistent restore operation of a VM.

20 Windows Server 2012 Support

Windows Server 2012 supports SnapManager for Hyper-V 1.2 onward. SnapDrive 6.5 for Windows is prerequisite software that must be installed on Windows Server 2012 to use SMHV. With SnapManager for Hyper-V 1.2, backup and restore of virtual machines will be supported only in SAN environments for Windows Server 2012.

20.1 Prerequisites

The prerequisites for SnapManager for Hyper-V 1.2 are as follows:

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

20.2 Feature Overview

SnapManager for Hyper-V 1.2 along with SnapDrive 6.5 for Windows will support all major SAN-based features in Windows Server 2012.

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

CSV 2.0 Support (CSVFS)

Note: In Windows Server 2012, 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, and this functions as a layer of abstraction above the NTFS file system for the storage volume. As a result, simultaneous reads/writes can be performed on the CSV LUN from different nodes. For more details on CSV 2.0, refer to http://technet.microsoft.com/en-us/library/jj612868.aspx. A CSV 2.0 volume will have two volume GUIDs:

- **NTFS volume GUID.** When a disk is created and partitioned with NTFS and before it is added to the CSV.
- **CSV volume GUID.** When a disk is added to the CSVs.

Best Practice

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

SnapManager for Hyper-V 1.2 supports virtual machines hosted on CSVFS volume type. The new CSVFS volume type has introduced a new CSV writer and CSV shadow copy provider. This has facilitated achieving distributed application-consistent backups. Section 20.3, “Asymmetric Clustering,” covers distributed application-consistent backup in detail.

20.3 Asymmetric Clustering

Asymmetric clustering is a feature with which users can create a shared disk or CSV among only a few nodes in a cluster.

Note: SnapManager for Hyper-V 1.2 does not support having virtual machines in such CSVs.

20.4 BitLocker Encryption

BitLocker was a data protection feature and was part of Windows 7 and Windows 2008 R2. This feature is now available with Windows Server 2012 with the additional functionality. The user will now be able to

encrypt cluster shared SAN volumes. For more information on BitLocker configuration, refer to http://technet.microsoft.com/en-us/library/hh831713.

SnapManager for Hyper-V will support BitLocker functionality for CSVs provisioned through SnapDrive 6.5 for Windows. Virtual machines can be hosted in encrypted CSVs.

20.5 New Virtual Hard Disk Format

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

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

SnapManager for Hyper-V 1.2 will support backup, restore, and replication of virtual machines in VHDx format.

Note: SnapDrive 6.5 for Windows currently cannot create LUNs beyond 16TB, and, therefore, NetApp advises creating a VHDx for sizes less than 16TB and to use other means of provisioning additional storage (pass-through disks, guest iSCSI initiator) on the VM.

20.6 Hyper-V Virtual Machine Live Migration

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

Best Practice
It is best to avoid SMHV-related operations within the virtual machine during live migration.

20.7 Hyper-V VM Storage Live Migration

This feature in Windows Server 2012 enables migrating virtual machine–related files to a different storage location without the VM having to undergo downtime. It is no longer necessary to take the virtual machine state offline when migrating to a different storage system.

Note: After migrating the virtual machine from one volume to another, restoring to a Snapshot copy taken in the earlier volume is not supported.

Best Practice
It is best to avoid SMHV-related operations during storage live migration. Otherwise, such operations could corrupt the virtual machine.

20.8 Windows Server 2012 Features Not Supported from SnapManager for Hyper-V 1.2 and SnapDrive 6.5 for Windows When Connected to NetApp Storage Systems Running in Clustered Data ONTAP Systems

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

- Hyper-V over SMB 3.0
- SMB over remote file shares
- SMB VSS for remote file shares (remote VSS)
- Virtual Fibre Channel
- Hyper-V replica

- Windows Server 2012 native thin provisioning
- Offload data transfer capability

21 SnapManager for Hyper-V 1.2 Backup Mechanism for Windows Server 2012

In Windows Server 2012, Microsoft introduced the **CSV proxy file system (CSVFS)**. The CSVFS provides a cluster shared storage LUN with a single and consistent file namespace while still using the underlying NTFS file system. In Windows Server 2012, the CSVs now appear as CSV file system, instead of NTFS (in Windows Server 2008 R2). For additional information on CSVFS architecture, refer to this link.

In Windows 2008 R2, CSV Hyper-V backup creates application-consistent backups on the each VM owner node. CSV ownership is moved to the VM owner node as part of the backup process. Hyper-V VSS writer then coordinates the freeze and thaw operations in the Hyper-V guest, and a subsequent hardware Snapshot is taken from the Hyper-V parent using the Data ONTAP VSS hardware provider (SnapDrive for Windows). This resulted in creation of a hardware Snapshot copy for each Windows cluster node, thereby introducing several scalability and space efficiency issues when the number of nodes in the cluster was increased.

In Windows Server 2012, CSVFS introduces “distributed application-consistent backups.” This allows backup of all the VMs in a cluster to be consistent in “one single application-consistent backup.” In order to achieve this distributed backup mechanism, Microsoft has introduced a new CSV writer and CSV provider.

- **CSV writer.** CSV writer serves the component-level metadata from the nonrequesting node for CSV volumes, and it functions as a proxy by including the Hyper-V writers from the remote node for the backup session.
- **CSV provider.** CSV provider coordinates the VSS back activities from all the Hyper-V writers on the partner cluster nodes to make the VM in an application-consistent state. Also, the CSV provider makes sure that CSV shadow copy volume is writable for the partner node Hyper-V writers during the auto recovery process.

Figure 28 illustrates SMHV 1.2 backup process for Windows Server 2012.

Figure 28) SMHV 1.2 backup process for Windows Server 2012.

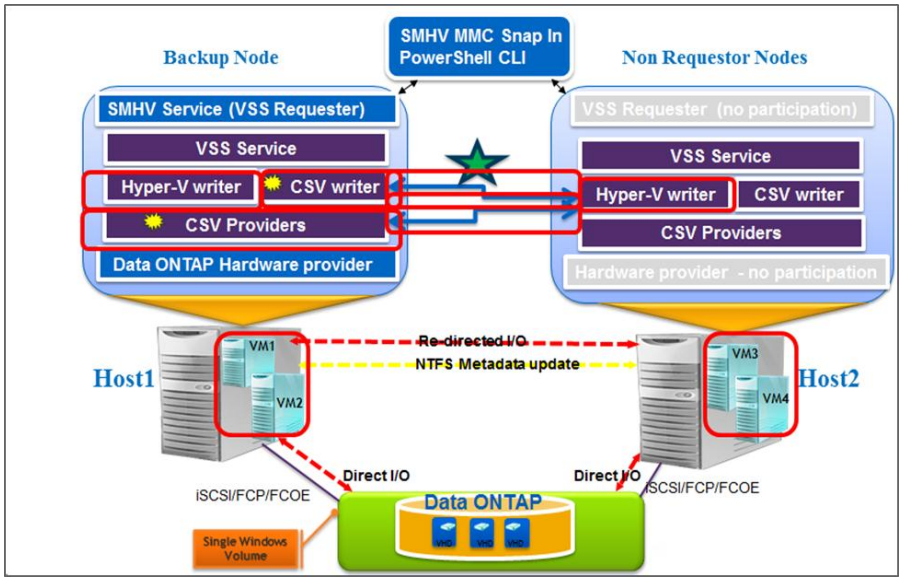

Initialization Phase

- The user initiates the backup operation from any node in the cluster using SMHV. SMHV redirects the backup operation to the Windows cluster owner node, which functions as a coordinator node throughout the entire backup operation.
- SMHV initializes the Microsoft VSS operation only in the coordinator node. This is unlike Windows Server 2008 R2, wherein VSS is initialized in each node of the Windows cluster, which is involved in the backup. This optimization improves the overall timing of the backup operation.
- SMHV gathers the metadata (files used by VMs) for all the VMs involved in the backup. Metadata for VMs that are local to the coordinator node is gathered by the Hyper-V writer running in the coordinator node.
- Metadata for the VMs that are not local to the coordinator node is gathered by the CSV writer running in the coordinator node. Internally the CSV writer in the coordinator node interacts with the Hyper-V writer in other nodes to get the metadata from all other nodes. So, unlike Windows Server 2008 R2, in which SMHV explicitly reaches out to each node to capture the metadata, this complication is handled by the new CSV writer in Windows 2012.

Prebackup Phase

- Hyper-V writer on the coordinator node quiesces the application writers inside the VM using the integration service.
- CSV software provider on the coordinator node interacts with the Hyper-V writers in all the other VM owner nodes to make sure that the state of the application running inside VM is consistent before starting the actual Snapshot copy of the volume.

Backup Phase

- VSS hardware provider on the coordinator node takes the backup Snapshot copy of the CSV volume.
- Hyper-V writer, by default, performs an autorecovery process on each VM owner node after the hardware Snapshot copy is created to remove any inflight transactions. Autorecovered changes are applied on the pseudo CSV Snapshot disk object exposed on the backup node, which is accessible from all the other VM nodes. This process makes the backups on each VM owner node application consistent with respect to CSV.

Postbackup Phase

- SMHV retrieves the VSS backup metadata and backup component documents and then modifies both the metadata to make it compatible with VSS required semantics.
- SMHV saves the backup metadata to the snapinfo folder.
- VSS Snapshot GUID is renamed to SMHV naming conventions.
- Applicable policy processing such as retention of older backups, SnapMirror updates, running any specified postscript, or generating ASUP™ notifications, is performed.

Note: Make sure that the “enable distributed backup” option Is checked in the backup dataset wizard.

Note: The distributed backup mechanism for Windows 2012 is not applicable for the crash-consistent backup feature in SMHV.

Note: It is recommended that all the VHD files belonging to a virtual machine are hosted on CSVFS LUNs only and not a mix of CSVFS and shared disks. This is because SMHV does not support such mixed-mode backups.

Best Practice

In order to achieve a successful backup and faster backup performance, it is recommended not to have more than 15 CSVFS LUNs in a single SMHV backup dataset that belong to the same NetApp storage system. In other words, virtual machines hosted on not more than 15 CSVFS LUNs belonging to the same storage system should be grouped together in a single dataset.

If we have 20 CSVFS LUNs hosted on a single NetApp storage system, it is recommended to create two datasets minimally and spread the virtual machines (CSVFS LUNs) evenly across these datasets.

To summarize, distributed application-consistent backups are faster since they avoid multiple backup requests to each node in the cluster. The entire backup operation is performed from the coordinator node (cluster owner) alone and by leveraging the new CSV writer and CSV shadow copy provider.

Also, distributed application-consistent backup is more space efficient since it creates only one Snapshot copy for each volume instead of creating one Snapshot copy for each node and volume combination. This space saving is huge if large numbers of nodes are involved in the backup. Also, Data ONTAP imposes a limit for the maximum number of Snapshot copies that could be stored for a volume, so considering that aspect, this enhancement would allow storing more backups for a VM.

22 Summary of SMHV Best Practices

The summary of SMHV best practices are provided in this section.

Best Practice

NetApp recommends to have one VM per LUN configured while deploying Hyper-V on a shared storage. All of the VHDs relative to a single VM (VM with multiple drives) can reside on single LUN provisioned as a shared storage to a WFC. It is a best practice for Windows 2008 Server R2 running Hyper-V deployed on standard shared storage volumes.

Best Practice

For SMHV, make sure that the following ports are kept open:

- 808: SnapDrive default port
- 4094: If SnapDrive is configured to use HTTP protocol
- 4095: If SnapDrive is configured to use HTTPS protocol

The default port number is 808. When SMHV is installed on a cluster, the same port number should be used across all nodes.

Best Practice

Having a SnapInfo LUN on a volume of its own is preferable.

Best Practice

When creating a dataset, you should select all VMs that reside on a particular Data ONTAP LUN. This enables you to get all backups in one Snapshot copy and reduce the space consumption on the storage system.

Best Practice

If you change a VM Snapshot copy file location to a different Data ONTAP LUN after creating the VM, you should create at least one VM Snapshot copy using Hyper-V manager before creating a backup using SMHV. If you change the Snapshot file location and do not create a VM Snapshot copy before creating a backup, the backup could fail.

Best Practice

The backup frequency, as well as the number of different backups performed against a dataset—for example, one backup running against dataset ds_1 weekly and another monthly—must be taken into account when specifying the retention policy so as not to exceed the maximum number of Snapshot copies per volume. Should the number of Snapshot copies exceed 255 on any given volume, future backups against that volume will fail.

Best Practice

Choose a backup retention level based on your backup creation and verification schedule. If a Snapshot copy deletion occurs, you should make sure that a minimum of one verified backup remains on the volume. Otherwise, you run a higher risk of not having a usable backup set from which to restore in case of a disaster.

Best Practice

For mission-critical VMs NetApp recommends enabling the “Allow Saved state VM backup” option.

Best Practice

SMHV cannot back up a VM that is actively undergoing migration. If a backup runs against a dataset that has VMs actively being migrated, an error is generated, and those particular VMs are not backed up. NetApp recommends that VMs be migrated only when a significant gain in performance can be achieved. This will improve not only the success rate of the backups, but the overall VM performance as well.

Best Practice

If the number of VHDs at the time of backup and restore is not same, the restored VM might have additional/fewer VHDs. If that is the case, NetApp recommends that the cluster configuration of the VM and its dependencies be manually updated.

Best Practices

Use active-active storage controller configuration to eliminate any SPOFs.

Use multipath HA with active-active storage configuration to get a better storage availability and higher performance.

More details on high-availability system configuration can be obtained from NetApp TR-3450: Active-Active Controller Overview and Best Practices Guidelines.

Best Practices

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

For Windows Server 2008 R2 servers, NetApp recommends Data ONTAP DSM 3.2R1 or higher.

For the currently supported multipathing software versions and related requirements, see the NetApp Interoperability Matrix.

23 SMHV Conclusion

SnapManager 1.0 for Hyper-V provides a rich feature set that allows IT organizations to take advantage of NetApp Snapshot and SnapMirror technologies to provide fast, space-efficient disk-based backups in a Hyper-V environment with NetApp storage while placing minimal overhead on the associated virtual infrastructure. The recommendations and examples in this report will help administrators get the most out of SMHV deployments.

Appendixes

Quick Steps to Deploy a Windows 2008 R2 Hyper-V Cluster Environment on NetApp Storage

Follow these steps to deploy a Windows 2008 R2 Hyper-V cluster environment on NetApp storage:

1. Install the NetApp storage system
   a. Create aggregates to support infrastructure.
   b. Create volumes to support clustered shared volume (CSV) infrastructure (turn on thin provisioning).
2. Perform server OS preparation
   a. Install Windows OS.
   – Hyper-V: role
   – Failover cluster feature

- NET 3.5 feature
- MPIO feature
- All patches

b. Install Microsoft hot fixes.

- KB975921
- KB974909
- KB975354 V2
- KB979743 V2
- KB974909 V2
- Install Windows Server 2008 R2 SP1
- KB2406705 V2
- KB978157

c. Install NetApp software.

- NetApp Windows Host Utility Kit 5.3
- NetApp MPIO 3.4
- SnapDrive 6.3 P2

3. Set up server network
   a. Network 1: Server management (ILO) (optional)
   b. Network 2: Client access (VM BRIDGE)
   c. Network 3: Live migration network (optional)
   d. Network 4: Heartbeat network
   e. Network 5: ISCSI network (as needed)
   f. Network 6: CSV network (for redirected I/O)(optional)
   g. HBA: FCP connections (as needed)
4. Set up SnapDrive for Windows
   a. Set up transport protocol defaults.
   b. Set up individual controllers.
   c. Provision the disks from the preferred controller IP address.
5. Set up cluster
   a. Create Windows cluster.
   b. Enable CSVs.
   c. Use SnapDrive to set up LUN to be used for quorum drive.
   d. Use Cluster Manager to set up failover cluster settings for quorum system.
6. Create CSVs
   a. Use SnapDrive for windows to create CSV.
   b. Open System Manager and convert LUNs to thin-provisioned LUNs.
7. Set up SnapManager for Hyper-V (SMHV)
   a. Use SnapDrive to create a single clustered drive to be used for the SnapInfo directory. Install SMHV on every node in the clusters.
   b. Install SMHV on every node in the clusters.
   c. Add cluster to the SMHV management console. (Refer to SMHV Installation and Administration Guide for help.)

d. Create a base dataset. (Refer to SMHV Installation and Administration Guide.)
e. Create a backup policy. (Refer to SMHV Installation and Administration Guide.)

8. Create VMs as needed
9. Open SMHV to add virtual machines to appropriate datasets
10. Repeat steps 8 and 9 as needed

How to Choose Your Hyper-V and VHD Storage Container Format

Customers have to make a choice when they need to decide what the appropriate storage container format is for deploying virtual machines using Hyper-V.

The summary in Table 16 is intended to make the decision-making process easier.

Table 16) Choosing Hyper-V and VHD storage container format.

Storage Container	Pros	Cons
Pass-through disk	• Fastest performance • Simplest storage path because file system on host is not involved • Better alignment under SAN • For shared storage based pass-through, no need to mount the file system on host and that might speed up VM live migration • Lower CPU use • Support very large disks	• VM snapshot cannot be created • Disk is being used exclusively and directly by a single virtual machine • Pass-through disks cannot be backed up by the Hyper-V VSS writer and any backup program that uses the Hyper-V VSS writer
Fixed size VHD	• Highest performance of all VHD types • Simplest VHD file format to give the best I/O alignment • More robust than dynamic or differencing VHD because of the lack of block allocation tables (redirection layer) • File-based storage container has more management advantages than pass-through disk • Expanding is available to increase the capacity of VHD • No risk of underlying volume running out of space during VM operations	• Up front space allocation might increase the storage cost when large number of fixed VHDs are deployed • Large fixed VHD creation is time-consuming • Shrinking the virtual capacity (reducing the virtual size) is not possible
Dynamically expanding or differencing VHD	• Good performance • Quicker to create than fixed-size VHD • Grow dynamically to save disk space and provide efficient storage usage • Smaller VHD file size makes it more nimble in terms of transporting across the network • Blocks of full zeros will not get allocated and thus save the space under certain circumstances • Compact operation is available to reduce the actual physical file size	• Interleaving of metadata and data blocks might cause I/O alignment issues • Write performance might suffer during VHD expanding • Dynamically expanding and differencing VHDs cannot exceed 2040GB • Might get VM paused or pull the VHD out if disk space is running out due to the dynamic growth • Shrinking the virtual capacity is

Storage Container	Pros	Cons
		not supported • Expanding is not available for differencing VHDs due to the inherent size limitation of parent disk • Defrag is not recommended because of inherent redirectional layer

SMHV: Virtual Machine Self-Management

If a VM belongs a host that has SMHV installed, and you install SMHV on that VM to use as a management console, you should not use SMHV to manage the host to which the VM belongs.

For example, if VM1 belongs to Host1 (with SMHV installed), and you install SMHV on VM1, you should not use SMHV to manage Host1 from VM1.

If you do this and try to restore the VM from itself, the VM will be deleted or restarted from Hyper-V Manager.

SMHV: Data ONTAP VSS Hardware Provider Requirement

Data ONTAP VSS hardware provider must be installed for SnapManager to function properly. Data ONTAP VSS hardware provider integrates the SnapDrive service and storage systems running Data ONTAP into the VSS framework. The Data ONTAP VSS hardware provider is now included with SnapDrive 6.0 or later and does not need to be installed separately.

Viewing Installed VSS Providers

To view the VSS providers installed on your host, complete these steps:

1. Select Start
2. Run and enter the following command to open a Windows command prompt: cmd.
3. At the prompt, enter the following command:

```
Vssadminlist providers
```

The output should be similar to the following:

```
Provider name: 'Data ONTAP VSS
Hardware Provider' Provider type:
Hardware
Provider Id: {ddd3d232-a96f-4ac5-8f7b-250fd91fd102}
Version: 6.2.0.xxxx
```

Verifying That the VSS Hardware Provider Was Used Successfully

To verify that the Data ONTAP VSS hardware provider was used successfully after a Snapshot copy was created, complete this step.

Navigate to System Tools > Event Viewer > Application in MMC and look for an event with the following values:

```
Source Event ID Description
The VSS provider has successfully completed CommitSnapshots for SnashotSetId id in n
milliseconds. Navsspr 4089
```

Note: VSS requires that the provider initiate a Snapshot copy within 10 seconds. If this time limit is exceeded, the Data ONTAP VSS hardware provider logs event ID 4364. This limit could be exceeded because of a transient problem. If this event is logged for a failed backup, retry the backup.

SMHV: When Virtual Machine Backups Take too Long to Complete

If a virtual machine contains several direct-attached iSCSI LUNs or pass-through LUNs, and SnapDrive for Windows is installed on the virtual machine, the virtual machine backup can take a long time. The Hyper-V writer takes a hardware snapshot of all the LUNs in the virtual machine using the SnapDrive for Windows VSS hardware provider. There is a Microsoft hotfix that uses the default system provider (software provider) in the virtual machine to make the snapshot. As a result, the Data ONTAP VSS hardware provider is not used for snapshot creation inside the child OS, and the backup speed increases. For more information on the Microsoft hotfix, see Knowledge Base article 975354 on the Microsoft support site at http://support.microsoft.com/.

SMHV: Redirected I/O and VM Design Considerations

While redirected I/O is handled in a Windows Server 2008R2 Hyper-V cluster, SMB API calls are made from one cluster node to the cluster and CSV owner. This involves metadata traffic and other SMB API calls that can affect performance significantly.

NetApp recommends that the user manually assign CSV and VM ownership to specific nodes in the cluster. SMHV backup datasets must be created and designed to back up all VMs in a single CSV owned by each specific node as follows:

1. Using SnapDrive for Windows, create one CSV per host cluster node, based upon tiers of storage as necessary. For example, create one CSV for fast SAS disk and one for SATA.
2. Using SCVMM, migrate VMs into their respective CSVs and assign ownership of those VMs to the same node that owns the CSV.

Note: All of VM migrations should be performed using SCVMM.

3. Create an SMHV dataset for each CSV and make sure that all VMs that reside in that CSV are placed into that dataset. For best results, do not allow VMs owned by multiple nodes to coreside within the same CSV.
4. Create a backup policy for each dataset that matches the customer's backup needs.
5. Using Failover Cluster Manager:
   a. Assign preferred ownership of each VM to its appropriate cluster node.
   b. Assign preferred ownership of each CSV to its appropriate cluster node.
   c. Before running each backup for each cluster node, assign cluster master ownership to the cluster node being backed up by that SMHV dataset. This is done through Failover Cluster Manager or using a Windows PowerShell script that can be executed by SMHV at the beginning of the backup job.

SMHV: Transferring Snapshot Copies to SnapVault or a Tape Device

In order to transfer SMHV Snapshot copies to SnapVault or a tape device, users can create a script and use the SMHV postscript feature in SMHV dataset policy option.

SMHV offers the following predefined variables, which the administrator can pass in order to achieve this:

- $VMSnapshot
- $SnapInfoName
- $SnapInfoSnapshot

During the post–policy execution phase, SMHV will replace the $VMSnapshot variable with the Snapshot name, $SnapInfoName with the time stamp of the backup, and $SnapInfoSnapshot with the snapinfo Snapshot name. You can access these variables from your scripts and do the necessary actions.

Here is a sample script that transfers SMHV Snapshot copies to a SnapVault system:

The following scripts are used:

- sv_update.ps1: Is used to update the Hyper-V VM snapshots to secondary storage
- sv_update_snapinfo.ps1: Is used to update SnapManager for Hyper-V to secondary storage
- update-vmsnapshot.bat: Batch file is used to start the sv_update.ps1 file with the correct parameters
- update-snapinfo.bat: Batch file is used to start the sv_update_snapinfo.ps1 file with the correct parameters

Prerequisites

1. Download and install the NetApp Data ONTAP PowerShell Tookit: https://communities.netapp.com/community/products_and_solutions/microsoft/powershell/data_ontap_powershell_toolkit_downloads
2. Unzip the PowerShell Toolkit to C:\Windows\System32\WindowsPowerShell\v1.0\Modules.
3. In Windows PowerShell, set the Set-ExecutionPolicy to "RemoteSigned"; otherwise, no scripts are allowed to be run. This needs to be done on every Hyper-V host.

Configuration Procedure for SnapVault Script

- Primary – test02
- Secondary – test03

1. Create SnapVault relationships:

```
snapvault start -S test02:/vol/csv01/csv01 test03:/vol/csv01/csv01
snapvault start -S test02:/vol/csv02/csv02 test03:/vol/csv02/csv02
snapvault start -S test02:/vol/csv03/csv03 test03:/vol/csv03/csv03
snapvault start -S test02:/vol/smhv_snapinfo/snapinfo test03:/vol/smhv_snapinfo/snapinfo
```

2. Set retention time on the secondary volumes:

   From test03:

```
snapvault snap sched csv01 sv_daily_testhv01 7@-
snapvault snap sched csv01 sv_daily_testhv02 7@-
snapvault snap sched csv01 sv_daily_testhv03 7@-
snapvault snap sched csv01 sv_daily_testhv04 7@-
snapvault snap sched csv01 sv_daily_testhv05 7@-
snapvault snap sched csv01 sv_daily_testhv06 7@-
snapvault snap sched csv01 sv_daily_testhv07 7@-

snapvault snap sched csv02 sv_daily_testhv01 7@-
snapvault snap sched csv02 sv_daily_testhv02 7@-
snapvault snap sched csv02 sv_daily_testhv03 7@-
snapvault snap sched csv02 sv_daily_testhv04 7@-
snapvault snap sched csv02 sv_daily_testhv05 7@-
snapvault snap sched csv02 sv_daily_testhv06 7@-
snapvault snap sched csv02 sv_daily_testhv07 7@-

snapvault snap sched csv03 sv_daily_testhv01 7@-
snapvault snap sched csv03 sv_daily_testhv02 7@-
snapvault snap sched csv03 sv_daily_testhv03 7@-
snapvault snap sched csv03 sv_daily_testhv04 7@-
snapvault snap sched csv03 sv_daily_testhv05 7@-
snapvault snap sched csv03 sv_daily_testhv06 7@-
snapvault snap sched csv03 sv_daily_testhv07 7@-

snapvault snap sched smhv_snapinfo  sv_daily 7@-
```

3. Enable SIS on the secondary volumes and start first dedupe run:

```
sis on /vol/csv01
sis on /vol/csv02
sis on /vol/csv03
sis on /vol/smhv_snapinfo

sis start -s /vol/csv01
sis start -s /vol/csv02
sis start -s /vol/csv03
sis start -s /vol/smhv_snapinfo
```

4. Create a "scripts" folder under C:\Program Files\NetApp\SnapManager for Hyper-V.

   This needs to be done on every Hyper-V host.

5. Place the Windows PowerShell scripts and batch files into the "C:\Program Files\NetApp\SnapManager for Hyper-V\scripts\" folder of every SMHV server.

   This needs to be done on every Hyper-V host.

6. Configure SMHV.
7. Create an SMHV dataset.
8. Create a policy for the dataset.
   a. Add the postscript "update-vmsnapshot.bat" (batch files that calls the Windows PowerShell script) to the policy.
   b. Also, add the parameters "`$VMSnapshot $SnapInfoName`" to the "Arguments" box.
9. Try the backup and review the result.
   a. Now, all the VM LUNs will have a backup with a consistent copy to SnapVault. Next step is to create a schedule for the snapinfo LUN.
10. Create a Windows task, which will kick off the snapinfo update script. Schedule this to run after the SMHV backup. Make sure it runs when the SMHV backup is finished. Otherwise, there will be no snapinfo Snapshot copy. The snapinfo Snapshot copy is only created after all the nodes of the SMHV cluster are finished with the backup.

Script 1: sv_update.ps1

```
if($ARGS.Length -lt 8)
{
   cls
   write $(" ");
   write $("Usage: sv_update.ps1 <primary_filer> <secondary_filer> <primary_volume>
<secondary_volume> <secondary_path> <retention_period> ");
   write $(" ");
   write $("Example: sv_update.ps1 filer01 filer02 vol1 sv_vol1 /vol/sv_vol1/qtree1 sv_daily ");
   write $(" ");
   exit(1);
}

$prifiler = $ARGS[0]
$secfiler = $ARGS[1]
$privol = $ARGS[2]
$secvol = $ARGS[3]
$secpath = $ARGS[4]
$sret = $ARGS[5]
$VMSnapshot = $ARGS[6]
$SnapInfoName = $ARGS[7]

#Reformat the VMSnapshot string and SnapINfo Snapshot string
$VMSnapshot_backup = $VMSnapshot+"_backup"

#Get the machine name of the HV host in lowercase for the retention period
$hvhost = $env:computername.ToLower()

#set retention period for this HV Host
```

```
$hvhostret = $sret+"_"+$hvhost

Import-Module dataontap


#check to see if the snapshot exist on the primary volume. If not we will exit the script.
Connect-NaController $prifiler
$str = Get-NaSnapshot -Targetname $privol -snapname $VMSnapshot
if ($str)
{
        # 1 - Initiates SnapVault transfer (update) from Secondary using last SnapManager for
Hyper-v snapshot.
        Connect-NaController $secfiler
        Start-NaSnapVaultSecTransfer $secpath -PrimarySnapshot $VMSnapshot

        # 2 - Simple time loop that will wait until SV update on Secondary (snapvault_secondary
volume) is done, before creating snapshot (Step 4).
        #     This script loops every 5 seconds until SnapVault status shows "Idle". Steps 5 & 6
should run once a month against Secondary.
        Connect-NaController $secfiler
        $var = $null
        while (!$var -or ($var.status -ne "Idle"))
        {
          $var = Get-NaSnapvaultSecStatus -Path $secpath
          start-sleep -seconds 5
        }

        # 4 - Initiates SnapVault transfer (update) from Secondary using last SnapManager for
Hyper-v snapshot.
        # This is the application persistent snapshot

        Connect-NaController $secfiler
        Start-NaSnapVaultSecTransfer $secpath -PrimarySnapshot $VMSnapshot_backup -
NoLunCloneExpansion 1

        # 5 - Simple time loop that will wait until SV update on Secondary (snapvault_secondary
volume) is done, before creating snapshot (Step 4).
        #     This script loops every 5 seconds until SnapVault status shows "Idle". Steps 5 & 6
should run once a month against Secondary.

        Connect-NaController $secfiler
        $var = $null
        while (!$var -or ($var.status -ne "Idle"))
        {
          $var = Get-NaSnapvaultSecStatus -Path $secpath
          start-sleep -seconds 5
        }

        # 4 - This archives (creates snapshot) on Secondary using the given retention schedule.
        Connect-NaController $secfiler
        Start-NaSnapvaultSecSnapshot -VolumeName $secvol -ScheduleName $hvhostret
exit(1);
}

else
{
  write $("Nothing to do there is no primary snapshot ");
exit(1)}
```

Script 2: sv_update_snapinfo.ps1

```
if($ARGS.Length -lt 6)
{
   cls
   write $(" ");
   write $("Usage: sv_update.ps1 <primary_filer> <secondary_filer> <primary_volume>
<secondary_volume> <secondary_path> <retention_period>");
   write $("Example: sv_update.ps1 filer01 filer02 vol1 sv_vol1 /vol/sv_vol1/qtree1 sv_daily");
   write $(" ");
   exit(1);
```

```
}

$prifiler = $ARGS[0]
$secfiler = $ARGS[1]
$privol = $ARGS[2]
$secvol = $ARGS[3]
$secpath = $ARGS[4]
$sret = $ARGS[5]


Import-Module dataontap
#check to see if the snapshot exist on the primary volume. If not we will exit the script.
Connect-NaController $prifiler
$str = Get-NaSnapshot -Targetname $privol -snapname $VMSnapshot
if ($str)
{
        # 1 - Pulls the last SnapManager for Hyper-v snapshot (using "smhv_snapinfo" key word)
from the snapvault_primary volume on the Primary.
        Connect-NaController $prifiler
        $LastSnapshot = get-nasnapshot $privol | ? { $_.Name -match "smhv_snapinfo" } | Sort-
Object AccessTimeDT -Descending | Select-Object -first 1

        # 2 - Initiates SnapVault transfer (update) from Secondary using last SnapManager for
Hyper-v snapshot.
        Connect-NaController $secfiler
        Start-NaSnapVaultSecTransfer $secpath -PrimarySnapshot $LastSnapshot.Name

        # 3 - Simple time loop that will wait until SV update on Secondary (snapvault_secondary
volume) is done, before creating snapshot (Step 4).
        #     This script loops every 5 seconds until SnapVault status shows "Idle". Steps 5 & 6
should run once a month against Secondary.
        Connect-NaController $secfiler
        $var = $null
        while (!$var -or ($var.status -ne "Idle"))
        {
          $var = Get-NaSnapvaultSecStatus -Path $secpath
          start-sleep -seconds 5
        }
        # 4 - This archives (creates snapshot) on Secondary using the given retention schedule.
        Connect-NaController $secfiler
        Start-NaSnapvaultSecSnapshot -VolumeName $secvol -ScheduleName $sret
exit(1);
}
else
{
  write $("Nothing to do there is no primary snapshot ");
exit(1)}
```

Script 3: update-vmsnapshot.bat (Batch File)

```
powershell -file "C:\Program Files\NetApp\SnapManager for Hyper-V\scripts"\sv_update_snapinfo.ps1
test02 test03 smhv_snapinfo smhv_snapinfo /vol/smhv_snapinfo/snapinfo sv_daily
```

Script 4: update-snapinfo.bat (Batch File)

```
powershell -file "C:\Program Files\NetApp\SnapManager for Hyper-V\scripts"\sv_update.ps1 test02
test03 csv01 csv01 /vol/csv01/csv01 sv_daily backup %1 %2
powershell -file "C:\Program Files\NetApp\SnapManager for Hyper-V\scripts"\sv_update.ps1 test02
test03 csv02 csv02 /vol/csv02/csv02 sv_daily backup %1 %2
powershell -file "C:\Program Files\NetApp\SnapManager for Hyper-V\scripts"\sv_update.ps1 test02
test03 csv03 csv03 /vol/csv03/csv03 sv_daily backup %1 %2
```

If the user does not want to use a postscript to send the Snapshot copy to the SnapVault or tape system, the user can use this script separately at a different time. In order to do this, we need to sort the Snapshot copies and select the latest Snapshot copy that needs to be sent to the secondary storage system. This can be done invoking the '`snap list`' command in Data ONTAP.

References

- Virtualization with Hyper-V: Supported Guest Operating Systems
  www.microsoft.com/windowsserver2008/en/us/hyperv-supported-guest-os.aspx
- Install the Hyper-V Role on a Full Installation of Windows Server 2008
  http://technet.microsoft.com/en-us/library/cc794929(WS.10).aspx
- NetApp TR-3701: NetApp and Microsoft Virtualization: Solution and Implementation Guide
  www.netapp.com/us/library/technical-reports/TR-3701.html
- Install the Hyper-V Role on a Server Core Installation of Windows Server 2008
  http://technet.microsoft.com/en-us/library/cc794852(WS.10).aspx
- Microsoft Hyper-V Server 2008 Configuration Guide
  www.microsoft.com/Downloads/details.aspx?familyid=E1E111C9-FA69-4B4D-8963-1DD87804C04F&displaylang=en
- Infrastructure Planning and Design Guide for System Center Virtual Machine Manager 2008 R2
  http://technet.microsoft.com/en-us/library/cc196387.aspx
- Nvspbind
  http://archive.msdn.microsoft.com/nvspbind
- VMM System Requirements
  http://technet.microsoft.com/en-us/library/cc764328.aspx
- Configuring a SAN Environment for VMM
  http://technet.microsoft.com/en-us/library/cc764269.aspx
- New Installation of VMM
  http://technet.microsoft.com/en-us/library/cc793149.aspx
- Configuring Virtual Networks
  http://technet.microsoft.com/en-us/library/cc816585(WS.10).aspx
- Hyper-V: Using Hyper-V and Failover Clustering
  http://technet.microsoft.com/en-us/library/cc732181(WS.10).aspx
- Hyper-V: Using Live Migration with Cluster Shared Volumes in Windows Server 2008 R2
  http://technet.microsoft.com/en-us/library/dd446679(WS.10).aspx
- Hyper-V Live Migration Overview and Architecture
  www.microsoft.com/downloads/details.aspx?FamilyID=FDD083C6-3FC7-470B-8569-7E6A19FB0FDF&displaylang=en
- Virtual Network Manager
  http://technet.microsoft.com/en-us/library/cc754263.aspx
- New in Hyper-V Windows Server 2008 R2 Part 1: Dedicated Networks
  http://blogs.technet.com/jhoward/archive/2009/05/04/new-in-hyper-v-windows-server-2008-r2- part-1-dedicated-networks.aspx
- NetApp Support
  http://support.netapp.com/
- NetApp Interoperability Matrix
  https://now.netapp.com/matrix/mtx/login.do
- High-Availability System Configuration
  www.netapp.com/us/products/platform-os/active-active.html
- NetApp TR-3450: Active-Active Controller Overview and Best Practices Guidelines
  www.netapp.com/us/library/technical-reports/tr-3450.html
- NetApp TR-3437: Storage Subsystem Resiliency Guide
  www.netapp.com/us/library/technical-reports/tr-3437.html
- NetApp Fibre Channel and iSCSI Configuration Guide
  http://now.netapp.com/NOW/knowledge/docs/san/fcp_iscsi_config/

- Remote LAN Management (RLM)
  http://now.netapp.com/NOW/download/tools/rlm_fw/info.shtml
- Windows Host Utilities 5.0 Installation and Setup Guide
  http://now.netapp.com/NOW/knowledge/docs/san/#windows
- Data ONTAP DSM for Windows MPIO Installation and Administration Guide
  http://now.netapp.com/NOW/knowledge/docs/san/#mpio
- Data ONTAP Network and File Access Management Guide
  http://now.netapp.com/NOW/knowledge/docs/ontap/
- NetApp SnapDrive for Windows
  www.netapp.com/us/products/management-software/snapdrive-windows.html
- NetApp SnapManager Family
  www.netapp.com/us/products/management-software/
- SnapManager for Hyper-V Installation and Administration Guide
  https://login.netapp.com/oamforms/login.html
- Performance Tuning Guidelines for Windows Server 2008 R2
  www.microsoft.com/whdc/system/sysperf/Perf_tun_srv-R2.mspx
- What's New in Windows Server 2008 R2 Hyper-V Performance and Scale?
  http://blogs.msdn.com/tvoellm/archive/2009/08/05/what-s-new-in-windows-server-2008-r2- hyper-v-performance-and-scale.aspx
- Microsoft Virtual Hard Disk (VHD) FAQ
  http://technet.microsoft.com/en-us/bb738381.aspx
- Virtual Hard Disk (VHD) Image Format Specification
  http://technet.microsoft.com/en-us/virtualserver/bb676673.aspx
- Performance Tuning Guidelines for Windows Server 2008 R2
  www.microsoft.com/whdc/system/sysperf/Perf_tun_srv.mspx
- Hyper-V and VHD Performance: Dynamic vs. Fixed
  http://blogs.technet.com/winserverperformance/archive/2008/09/19/hyper-v-and-vhd- performance-dynamic-vs-fixed.aspx
- Data ONTAP Block Access Management Guide for iSCSI or FC
  http://now.netapp.com/NOW/knowledge/docs/ontap/rel731_vs/pdfs/ontap/bsag.pdf
- Data ONTAP Commands Manual Page Reference, Volumes 1 and 2
  https://now.netapp.com/AskNOW/search?action=search&search_collections=kbase&search_collections=bugs&search_collections=docs&search_collections=tools&query=Data%20ONTAP%20Commands%20Manual%20Page%20Reference
- Planning for Disks and Storage
  http://technet.microsoft.com/en-us/library/dd183729(WS.10).aspx
- Configuring Disks and Storage
  http://technet.microsoft.com/en-us/library/ee344823(WS.10).aspx
- NetApp TR-3505: NetApp Deduplication for FAS: Deployment and Implementation Guide
  www.netapp.com/us/library/technical-reports/tr-3505.html
- Microsoft KB 302577: http://support.microsoft.com/kb/302577
- Microsoft KB 958184: http://support.microsoft.com/kb/958184
- Microsoft KB 961804: http://support.microsoft.com/kb/961804/en-us
- SnapMirror Sync and SnapMirror Semi-Sync Overview and Design Considerations
  www.netapp.com/us/library/technical-reports/tr-3326.html
- SnapMirror Async Overview and Best Practices Guide
  www.netapp.com/us/library/technical-reports/tr-3446.html

- Configuring Alarms
  http://now.netapp.com/NOW/knowledge/docs/DFM_win/rel36r1/html/software/opsmgr/monitor5.htm
- Managing Aggregate Capacity
  http://now.netapp.com/NOW/knowledge/docs/DFM_win/rel36r1/html/software/opsmgr/filesys4.htm
- Operations Manager
  www.netapp.com/us/products/management-software/operations-manager.html
- A Description of the Diskpart Command-Line Utility
  http://support.microsoft.com/default.aspx?scid=kb;en-us;300415
- GNU ext2resize
  http://ext2resize.sourceforge.net/
- NetApp Data ONTAP PowerShell Toolkit
  http://communities.netapp.com/docs/DOC-6162#
- Data ONTAP Block Access Management Guide
  http://now.netapp.com/NOW/knowledge/docs/ontap/rel733/pdfs/ontap/bsag.pdf
- SnapDrive for Windows Installation and Administration Guide
  http://now.netapp.com/NOW/knowledge/docs/snapdrive/relsnap63/pdfs/admin.pdf
- Data ONTAP 7.3 System Administration Guide
  http://now.netapp.com/NOW/knowledge/docs/ontap/rel733/pdfs/ontap/sysadmin.pdf

Knowledge Base Articles

- KB ID: 3011206: SMHV: Can SnapManager 1.0 for Hyper-V exclude Virtual Hard Disks from backups?
- KB ID: 1010146: SMHV: How to manually restore a Hyper-V virtual machine from a Snapshot backup
- KB ID: 1011587: How to migrate a Hyper-V VM to support SnapManager for Hyper-V Backup
- KB ID: 2010899: SMHV: Backups fail for Hyper-V Virtual Machines containing Passthru or iSCSI in Guest Disks
- KB ID: 1010887: SMHV: How to setup SnapInfo Logical Unit Number (LUN)
- KB ID: 2010607: SMHV: Creation of two snapshots for every backup
- KB ID: 2014905: SnapManager for Hyper-V backups fail to complete even though all Virtual Machines are located on NetApp LUNs
- KB ID: 2014900: SnapManager for Hyper-V backup sets that contain Windows XP fail
- KB ID: 2014928: SMHV: During backup of CSV, hosts report `NO_DIRECT_IO_DUE_TO_FAILURE`.
- KB ID: 2014933: SMHV: Cluster Shared Volume goes offline after backup
- KB ID: 2639032: "0x0000003B," "0x00000027," and "0x0000007e" Stop errors when a connection to a CSV is lost on a Windows Server 2008 R2-based failover cluster
- KB2517329: Performance decreases in Windows Server 2008 R2 when the Hyper-V role is installed on a computer that uses Intel Westmere or Sandy Bridge processors
- KB2552040: A Windows Server 2008 R2 failover cluster loses quorum when an asymmetric communication failure occurs
- KB2522766: The MPIO driver fails over all paths incorrectly when a transient single failure occurs in Windows Server 2008 or in Windows Server 2008 R2
- KB2528357: Nonpaged pool leak when you disable and enable some storage controllers in Windows 7 or in Windows Server 2008 R2
- KB2770917 – This is a Windows Server 2012 KB fix.

  This is to fix the following error:
  "Error: Vss Requestor - Backup Components failed. Writer Microsoft Hyper-V VSS Writer involved in backup or restore encountered a retryable error. Writer returned failure code 0x800423f3. Writer state

is 8." This issue is caused by inclusion of direct attached iSCSI LUNs or pass-through disks in the VSS backups.

Version History

Version	Date	Document Version History
1.0	June 2009	Initial release.
2.0	January 2010	Updates to Hyper-V storage options, remove duplicate content with and add links to NetApp TR-3701.
3.0	January 2011	• In addition to minor formatting changes, the following areas of content were added and/or updated: • NetApp FlexClone volumes, NetApp Snapshot copies. And Virtual Machine Provisioning. File System Alignment and LUN Types. • Backup and Recovery using NetApp Snapshot copies, including many scripts in the Appendix for Backup and Restore of VMs using SnapDrive and VSS/Diskshadow. Disaster Recovery and High Availability using NetApp SnapMirror, including scripts in the Appendix for replication and recovery of infrastructure using NetApp SnapMirror. • Reformatted for publishing through book. Reorganization of outline and Table of Contents. Removed example scripts and added to NetApp Communities Site for Data ONTAP PowerShell Toolkit. • Removed content related to Microsoft Virtual Disk Service (VDS) while configuring SCVMM R2.
4.0	June 2012	Consolidation and updates of TR-3702, "NetApp Storage Best Practices for Microsoft Virtualization," and TR-3805, "SnapManager 1.0 for Hyper-V Best Practices."
5.0	January 2013	Updated with Windows Server 2012 support.

Acknowledgements

The author would like to thank the following people for their contributions:

- Anagha Barve, Member of Technical Staff, Microsoft Business Unit
- Atul Bhalodia, Senior Engineer, Microsoft Business Unit
- Chance Bingen, Escalation Engineer
- Chris A. Collins, Enterprise Infrastructure Architect
- John Fullbright, Reference Architect, Microsoft Business Unit
- Vineeth Karinta, Member of Technical Staff, Microsoft Business Unit

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

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

Go further, faster®

NetApp

www.netapp.com