

Using EMC Celerra IP Storage with Windows Server 2008 R2 Hyper-V over iSCSI

Best Practices Planning

Abstract

This white paper provides a detailed overview of EMC[®] Celerra[®] Network Server with Microsoft Hyper-V R2, the bare-metal hypervisor that is built into the 64-bit Windows Server 2008 R2. It includes added features such as Live Migration and Cluster Shared Volume (CSV) that enhance the availability and manageability of virtual machines. This white paper covers the use of Celerra protocols and features in support of the Hyper-V R2 virtualization environment.

August 2009

Copyright © 2009 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part Number h6530

Table of Contents

Executive summary	5
Introduction	5
Audience	5
Terminology	5
Celerra overview	7
Data Mover.....	8
Control Station	8
Basics of multiprotocol support on Celerra	8
Native iSCSI.....	9
CIFS and NFS.....	9
Microsoft Hyper-V	9
Live migration and quick migration	11
CSV	11
Hyper-V storage features.....	11
VHD	11
Passthrough disk	12
Dynamic virtual machine storage	12
Failover cluster storage	12
Hyper-V network concepts.....	13
Management console	13
IP storage	13
Virtual machine networks	14
Celerra storage provisioning for Hyper-V.....	17
Naming of storage objects	17
Celerra iSCSI configuration for Hyper-V.....	18
LUN configuration options.....	23
VHD	23
Passthrough disk.....	24
Microsoft iSCSI Software Initiator within a virtual machine.....	25
Virtual provisioning	26
File system virtual provisioning	26
iSCSI virtual provisioning	28
Partition alignment	29
iSCSI LUN extension	29
Celerra iSCSI LUN snapshot.....	31
Test/Dev environment.....	31
Backup/Restore data stored in a virtual machine	32
Snapping iSCSI LUN that is connected through the Microsoft iSCSI initiator within a virtual machine.....	33
Disk signature resolution.....	33
Replication	34
Conclusion	34

References	35
-------------------------	-----------

Executive summary

Virtualization is a familiar term used in technical conversations for many years. Administrators have become accustomed to almost all things virtual, from virtual memory to virtual networks to virtual storage.

In 2008, Microsoft released Hyper-V, its first bare-metal hypervisor-based technology, built into Windows Server 2008. The goal of Hyper-V was to provide a framework for companies to effectively partition and use their hardware resources and at the same time, compete in the virtual server space with products like VMware ESX and Citrix XenServer. With the Hyper-V architecture, hardware-assisted, x64-based systems can run independent, virtual environments with different operating systems and resource requirements within the same physical server.

In 2009, the release of Windows Server 2008 R2 introduced two anticipated features for Hyper-V: Live Migration and Cluster Shared Volume (CSV). These features work in a Windows failover clustering environment and additionally, can leverage EMC® Celerra® Network Server, which offers iSCSI logical unit numbers (LUNs) as a storage option to create virtual machines and virtual disks. The inclusion of the Microsoft iSCSI Software Initiator in Windows Server 2008 provides ubiquitous SAN connectivity for customers using existing Ethernet infrastructure. The iSCSI connectivity in Windows Server 2008 combined with the Hyper-V role offers unique benefits for virtualized environments. Celerra also provides advanced scalability and reliability by combining some of the benefits of Fibre Channel with the flexibility and ease of use of IP storage.

Introduction

Several Microsoft products/features such as System Center Virtual Machine Manager (SCVMM), System Center Operations Manager (SCOM), Performance and Resource Optimization (PRO), and Data Protection Manager (DPM) offer integrated solutions to manage virtual environments running Hyper-V. However, this white paper focuses specifically on Windows Server 2008 R2 with Hyper-V because Windows is the platform that supports Celerra IP storage.

Microsoft Hyper-V environments can take full advantage of IP block storage by using the iSCSI protocol. This provides a method to integrate virtualized computing to virtualized storage, offering a dynamic set of capabilities within the data center, and resulting in improved performance and system reliability. This white paper describes how to effectively use EMC Celerra Network Server in a Windows Server 2008 R2 with Hyper-V environment.

Audience

System administrators, system architects, and anyone who is interested in how Hyper-V can be integrated with iSCSI storage can benefit from this white paper. The reader who is familiar with server virtualization products (such as Microsoft Virtual Server) and EMC Celerra can benefit from this white paper and can easily understand the concepts that are presented.

Terminology

- **child partition:** Partition created by the hypervisor in response to a request from the parent or root partition. There are a couple of key differences between a child partition and a parent partition. Child partitions are unable to create new partitions. Child partitions do not have direct access to devices (any attempt to interact with hardware directly is routed to the parent partition) or memory (when a child partition tries to access the memory, the hypervisor/virtualization stack remaps the request to different memory locations).
- **Common Internet File System (CIFS):** File-sharing protocol based on the Microsoft Server Message Block (SMB). It allows users to share file systems over the Internet and intranets. Microsoft introduced a new version of the SMB protocol (SMB 2.0 or SMB2) with Windows Vista in 2006.

-
- **Cluster Shared Volume (CSV):** Windows Server 2008 R2 feature that enables multiple cluster nodes to access a shared LUN simultaneously, so that clustered virtual machines with virtual hard disk (VHD) files stored on the shared LUN can fail over independently.
 - **Data Execution Prevention (DEP):** Security feature intended to prevent an application or service from executing code from a non-executable memory region.
 - **differencing disk:** Type of VHD file that is associated with another disk that you want to leave intact in a parent-child relationship. You can make changes to the data or operating system without affecting the parent disk. Therefore, you can easily roll back the changes, if required.
 - **dynamically expanding disk:** Type of VHD file that grows, up to the size specified, as data is stored to the disk. The file does not shrink automatically when data is deleted.
 - **guest operating system:** Operating system that runs within a virtual machine or child partition.
 - **hypervisor:** Virtualization platform that enables multiple virtual machines to run in isolation on a physical host at the same time. Also known as a virtual machine monitor.
 - **Hyper-V:** Next-generation hypervisor-based server virtualization technology in Windows Server 2008. Its current guest operating system (OS) support includes Windows and Linux.
 - **Hyper-V Manager:** Microsoft Management Console (MMC)-based management interface that enables administrators to perform basic operations such as provisioning or managing virtual machines. It is installed when a Hyper-V server role is enabled on the Windows Server 2008 R2 full installation. The Windows Server Core version does not have a GUI interface.
 - **iSCSI target:** iSCSI endpoint, identified by a unique iSCSI name, which executes commands issued by the iSCSI initiator.
 - **Live Migration:** High-availability feature in Windows Server 2008 R2 that enables a running virtual machine to migrate from one physical host to another seamlessly.
 - **logical unit number (LUN):** Identifying number of a SCSI or iSCSI object that processes SCSI commands. The LUN is the last part of the SCSI address for a SCSI object. The LUN is an ID for the logical unit, but the term is often used to refer to the logical unit itself.
 - **Network File System (NFS):** Distributed file system providing transparent access to remote file systems. NFS allows all network systems to share a single copy of a directory.
 - **parent partition:** Partition that is capable of calling the hypervisor and requesting that new partitions be created. In the first release of Hyper-V, the parent and root partitions were the same and only one parent partition was possible.
 - **partition:** Basic entity managed by the hypervisor. It is an abstract container that consists of isolated processor and memory resources with policies on device access. A partition is a lighter weight concept than a virtual machine and can be used outside of the context of virtual machines to provide a highly isolated execution environment.
 - **passthrough disk:** Physical disks/LUNs that are directly attached to a virtual machine. A passthrough disk “bypasses” the file system layout at the parent partition. It provides a disk performance that is on par with a physical disk, but it does not support some of the functionality of virtual disks, such as virtual machine snapshots.
 - **root partition:** First partition on the computer. Specifically, this is the partition that is responsible for initially starting the hypervisor. It is also the only partition that has direct access to memory and devices.
 - **System Center Virtual Machine Manager (SCVMM):** Member of the System Center suite of server management products that provides unified management of physical and virtual machines, consolidation of underused physical servers, and rapid provisioning of new virtual machines by leveraging the expertise and investments in the Microsoft Windows Server technology. SCVMM 2008 supports Hyper-V and provides compatibility with VMware Virtual Infrastructure 3 through VirtualCenter integration.
 - **Virtual Hard Disk (VHD):** Common virtualization file format that captures the entire virtual machine operating system and application stack in a single file stored on a file system in the parent partition.

- **Virtual local area network (VLAN):** Group of devices physically residing on different network segments, but communicating as if they resided on the same network segment. VLANs are configured by management software and are based on logical connections as compared to physical connections for increased administrative flexibility.
- **virtual machine:** Virtualized x86 or x64 PC on which a guest operating system and an associated application run.
- **virtual machine configuration file:** XML file containing a virtual machine configuration that is created by the New Virtual Machine Wizard. It contains a virtual machine ID that uniquely identifies each virtual machine.
- **virtually provisioned iSCSI LUN:** iSCSI LUN without reserved space on the file system. Adequate file system space must be available for allocation whenever data is added to the LUN.

Celerra overview

The Celerra network-attached storage (NAS) product family covers a broad range of configurations and capabilities that scale from the midrange to the high end of networked storage. Although differences exist along the product line, there are some common building blocks. These basic building blocks are combined to encompass a broad, scalable product line with consistent support and configuration options.

A Celerra frame provides n+1 power and cooling redundancy and supports a scalable number of physical disks, based on the model you select and the needs of your solution. For this white paper, there are two additional building blocks:

- Data Mover
- Control Station

Data Movers move data back and forth between the local area network (LAN) and the back-end storage (disks). The Control Station is the management station for the system. Celerra is configured and controlled through the Control Station. [Figure 1](#) shows how Celerra works.

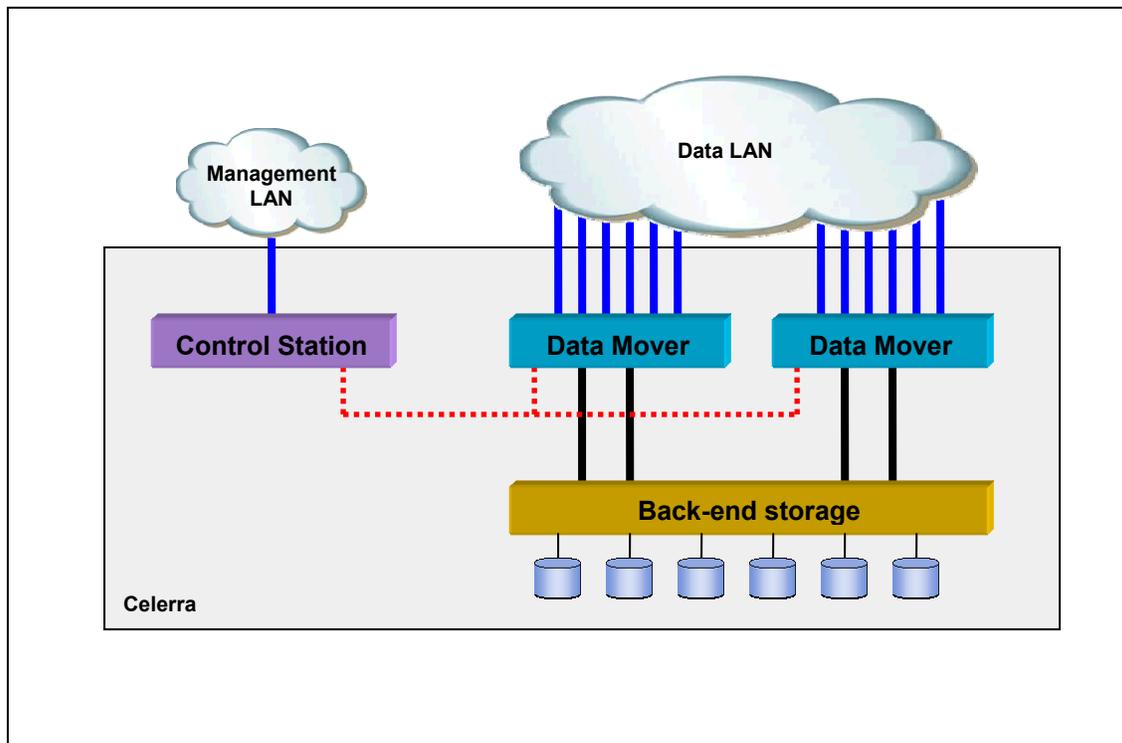


Figure 1 Celerra block diagram

Data Mover

Celerra has one or more Data Movers installed in its frame. You can think of a Data Mover as an independent server that is running EMC's optimized NAS operating system, data access in real time (DART). Each Data Mover has its own network ports and network identity, and it also has its own connections to the back-end storage. In many ways, a Data Mover operates as an independent server, bridging the LAN and the back-end storage disk array.

Multiple Data Movers are grouped together as a single system for high availability and ease of use. For high availability, Celerra supports a configuration in which one Data Mover acts as a standby for one or more active Data Movers. When an active Data Mover fails, you can start the standby quickly to take over the identity and storage of the failed device. For ease of use, logically group all Data Movers in a cabinet for administration as a single system through the Control Station.

Control Station

The Control Station is the single point of control and management of a Celerra frame. Regardless of the number of Data Movers or disk drives in the system, the administration of the system is done through the Control Station. The Control Station provides the interface to configure the Data Movers and back-end storage, and it also provides heartbeat monitoring of the Data Movers. It is important to note that if a Control Station is inoperable for any reason, the Data Movers continue to operate normally. However, the Celerra architecture provides an option for a redundant Control Station to support the continuous management of Celerra for an increased level of availability.

The Control Station runs a version of the Linux OS that EMC has optimized for Celerra and NAS administration. [Figure 1](#) on page 7 shows a Celerra system with two Data Movers. Based on the product model, the Celerra NAS family supports a maximum of 14 Data Movers.

Basics of multiprotocol support on Celerra

Celerra provides access to block and file data by using iSCSI, Common Internet File System (CIFS), and Network File System (NFS). [Figure 2](#) shows the storage protocols that provide standard TCP/IP network services.

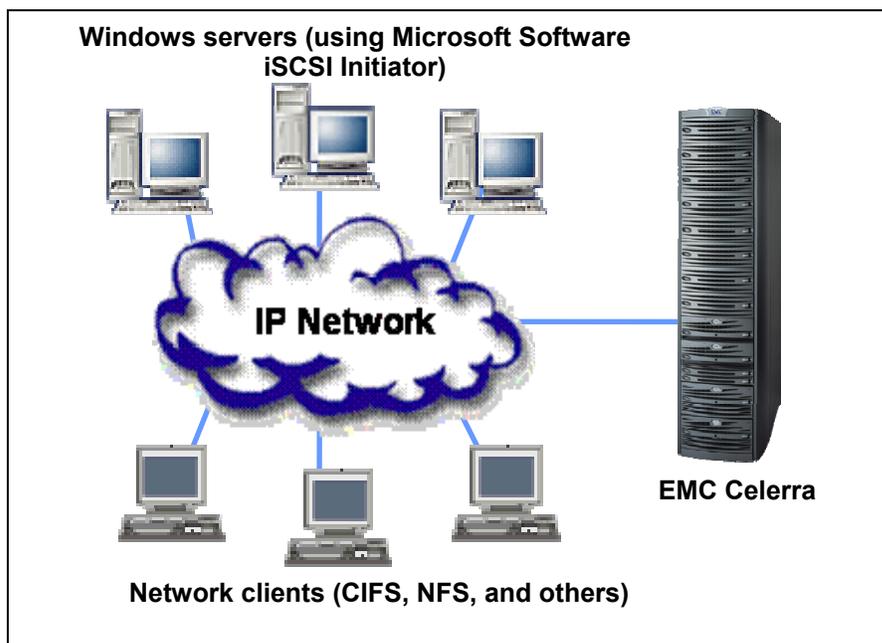


Figure 2 Celerra iSCSI topology

Native iSCSI

Celerra delivers a native iSCSI solution as an iSCSI target to which multiple hosts can access. This means that when deployed, all connected devices have TCP/IP interfaces capable of supporting the iSCSI protocol. A simple view of this topology is shown in Figure 3. No special hardware, software, or additional configuration is required. The administrator only needs to configure the two endpoints: the iSCSI target on Celerra and the iSCSI host initiator on the server. This option is easier to configure and manage than a typical iSCSI bridging deployment.

A bridging solution requires additional hardware that acts as the bridge between the Fibre Channel network connected to the storage array and the TCP/IP network connected to the host initiator. Bridging solutions were useful in the early days of iSCSI before the general adoption of native iSCSI targets. However, the addition of the bridge adds to the hardware cost and configuration complexity. This is not necessary when deploying iSCSI on Celerra. With a bridging deployment, you must configure the iSCSI target on the storage array, the iSCSI host initiator on the server, and the corresponding connections on the bridge.

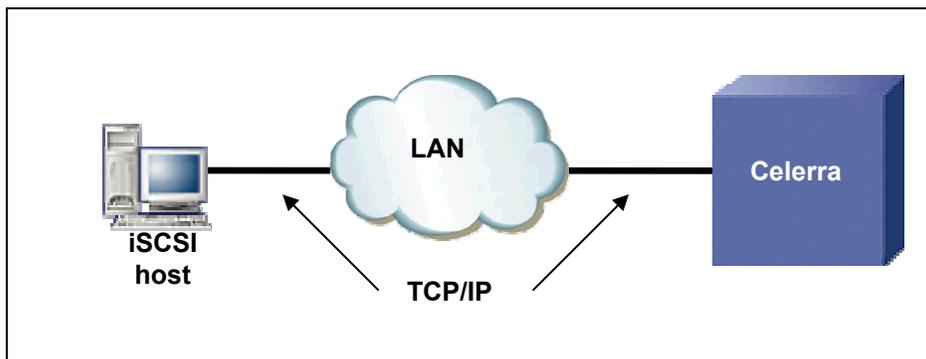


Figure 3 Native iSCSI topology

CIFS and NFS

CIFS and NFS are remote file-sharing protocols offered by the Celerra network storage system. Celerra provides support for CIFS and NFS servers through one or more Data Movers, thus enabling CIFS and NFS clients to safely access shared storage across the network. Celerra provides added functionality to the CIFS and NFS services by offering support for snapshot copies in addition to integrated replication and backup technologies. The core component that provides the CIFS and NFS services is the Celerra Data Mover.

Microsoft Hyper-V

Microsoft Hyper-V is a virtualization software that provides server consolidation by enabling several instances of similar and dissimilar operating systems to run as virtual machines on one physical machine. This cost-effective, highly scalable virtual machine platform offers advanced resource management capabilities. Hyper-V minimizes the total cost of ownership (TCO) of computing infrastructure by:

- Increasing resource utilization
- Decreasing the number of servers and all associated costs
- Maximizing server manageability

There are two requirements on the hardware platform:

- The server must support and enable hardware-assisted virtualization technology such as Intel Virtualization Technology (VT) or AMD Virtualization (AMD-V).
- The hardware must support and enable Data Execution Prevention (DEP). This means enabling the execute disable (XD) bit on the Intel platform or the no execute (NX) bit on AMD.

The Hyper-V feature is enabled as a server role in the x64 version of the Windows Server 2008 R2 Standard, Enterprise, and Datacenter editions. When Hyper-V is enabled, the Windows Server 2008 R2 environment becomes the parent partition, which is where the guest OS is created, device drivers are installed and updated, and the management interface is accessed. The guest OS is created in the context of a child partition, which interacts with the parent partition by using the hypercall application program interface through the VMBus. VMBus is the communication channel between the virtual service provider (VSP) in the parent and the virtual service client (VSC) in the child.

Before loading the parent and child partitions, the Hyper-V hypervisor is loaded directly on top of the hardware, enabling virtual machines to run above the virtualization layer provided by the hypervisor. Figure 4 shows the relationship between each component in an architectural diagram.

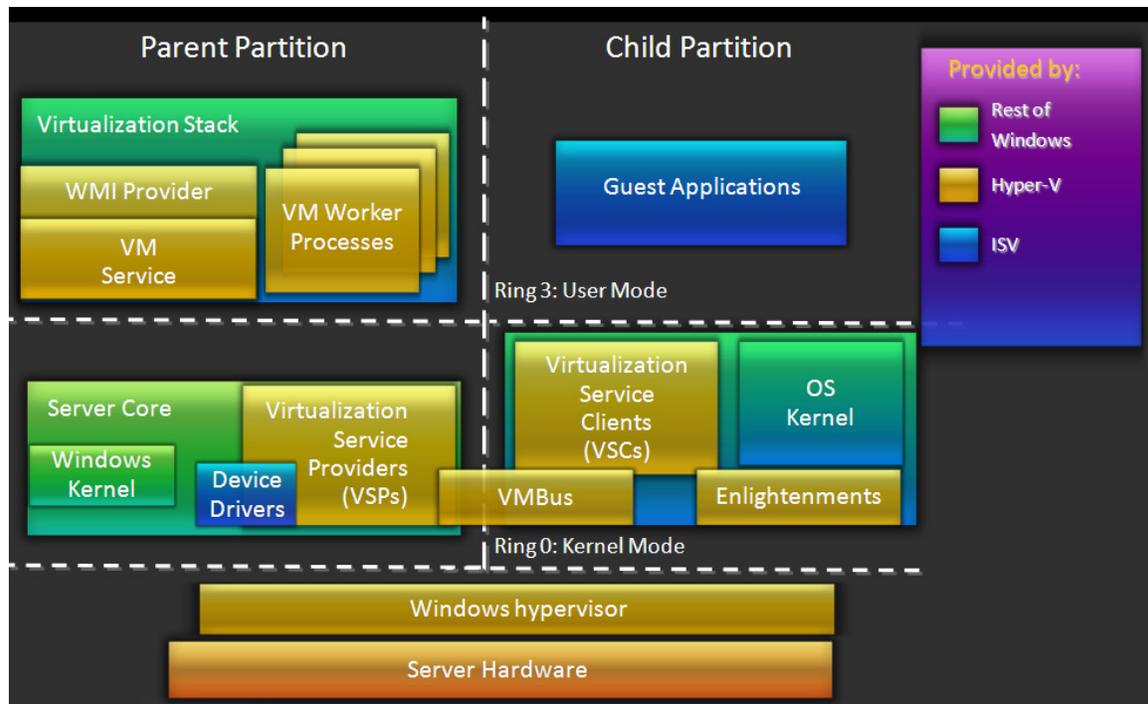


Figure 4 Architecture of Microsoft Hyper-V

The parent partition runs in one of the two variations of Windows Server 2008 R2: the full installation Server or Server Core. Server Core is a minimal server installation option for computers running on the Windows Server 2008 operating system. Server Core provides a low-maintenance server environment with a limited set of components, which are mandatory for core OS functions. The focus on required components reduces the attack surface and the amount of code that would require patching. Server Core is managed from the command line or from a remote server.

Note: Windows Server 2008 R2 Server Core installation includes the Microsoft iSCSI software initiator GUI (iSCSICPL). You can launch the GUI from the command line.

Other than the Windows Server 2008 R2 versions of Hyper-V, Microsoft offers a free, standalone virtualization edition that is known as Microsoft Hyper-V R2 Server. Similar to Server Core, it is a stripped-down version that includes only a CLI with limited capabilities.

Hyper-V environment is managed by using the Hyper-V Manager GUI that comes with Windows Server 2008 R2. For Server Core or Hyper-V Server, there is no GUI support locally. Therefore, you need to start the Hyper-V Manager remotely on another management server.

In addition to the Hyper-V Manager that performs basic management tasks, you can perform advanced features such as live or quick migration and virtual machine cloning by using an optional Microsoft

management product, SCVMM. This enables an administrator to centrally manage a number of Hyper-V servers in addition to a library of virtualization objects. SCVMM is a product that you need to purchase and license separately.

Live migration and quick migration

Prior to Windows Server 2008 R2, it was possible to migrate a running virtual machine from one physical host to another through quick migration in Windows Server 2008. Quick migration incurs a minimal downtime when the virtual machine is transitioned to another host. Users that connect to the virtual machine may experience a brief interruption during the quick migration and may have to re-establish the disconnected sessions.

With the introduction of live migration in Windows Server 2008 R2, the transition of the virtual machine migration is so seamless that users who connect to the virtual machine do not experience noticeable service interruptions. Windows Server 2008 R2 supports both live and quick migration options. You can use quick migration to move multiple virtual machines simultaneously. However, the number of concurrent live migration operations is limited among cluster nodes within a failover cluster. On a server running Hyper-V, only one live migration (to or from the server) can be in progress at a given time. There is also a list of prerequisites for completing live migration successfully. For instance, a live migration can fail if there is inadequate network throughput between the source and destination nodes. The Windows Server 2008 R2 documentation provides the complete list of requirements, limitations, and technical details of live and quick migrations.

CSV

CSV, along with live migration, is the other prominent new feature added to Windows Server 2008 R2 with respect to Hyper-V enhancements. You must have at least one NTFS file system on a shared disk to convert it to a CSV. When a CSV is used in a failover cluster, multiple cluster-enabled virtual machines running on different nodes can simultaneously access their virtual hard disks (VHD), which is stored on a shared LUN, so that each virtual machine can fail over independently.

On the contrary, if a traditional or non-CSV shared volume is used in a failover cluster, one node has exclusive ownership and access to the shared LUN. If multiple cluster-enabled virtual machines have their VHDs stored on the shared LUN, they all have to fail over together along with the disk resources.

Live migration can work with or without a CSV. When a CSV is used, it is accessed by each cluster node through a common namespace that typically looks like `C:\ClusterStorage\VolumeX` where X is a numeric value that represents a CSV volume number.

CSV comes with a feature known as dynamic I/O redirection. If a node loses storage connectivity to a CSV, I/Os are rerouted to another cluster node through the network interconnect. Windows Server 2008 R2 documentation provides more information.

Hyper-V storage features

There are two ways to provision storage for virtual machines: virtual hard disk (VHD) or passthrough disk. This section also covers dynamic virtual machine storage and failover cluster storage.

VHD

VHD is a disk image format implemented by Microsoft. Its specification is available to third parties so that other hypervisors such as XenServer can leverage this technology. VHDs are created as .vhd files that reside on the native host file system of the server running Hyper-V and support a wide range of storage types including iSCSI, or theoretically, any other file systems recognized and accessible by Windows Server 2008 R2.

You can create three types of VHDs with Hyper-V:

-
- **Dynamically expanding** — This is thin provisioning where the .vhd file grows, up to the size that you specify, as data is stored to the disk. The .vhd file does not shrink automatically when data is deleted.
 - **Fixed size** — The .vhd file is preallocated to the amount of space you specify for the disk size, regardless of how much data is saved to the virtual hard disk.
 - **Differencing** — This type of disk is associated in a parent-child relationship with another disk that you want to leave intact. You can make changes to the data or operating system without affecting the parent disk. Therefore, it is easy to roll back the changes.

Based on the disk type, Microsoft provides a **Edit Virtual Hard Disk Wizard** utility that can perform a variety of maintenance operations.

You can compact, convert, and expand a dynamically expanding disk. Compact shrinks a .vhd file by removing blank spaces that remain when the data is deleted from the disk. Convert causes a dynamic disk to become a fixed-size disk by copying the contents. Expand expands the capacity of the virtual disk.

You can convert a fixed-size disk to a dynamic or expanded disk only.

You can compact or merge a differencing disk where the changes stored in a differencing disk are merged into the parent disk or into some other disk.

You can inspect a virtual disk (select Inspect Disk in the Actions pane in Hyper-V Manager) to determine its type, location, and actual or maximum size. Because a differencing disk depends on its parent disk, if the filename of the parent disk is renamed, it cannot locate the parent resulting in a broken chain. An option is available to reconnect the chain by manually identifying the correct parent. Unless you are careful about the way you name the VHD files, it can turn into a nightmare to locate the right parent disk out of a multitude of disks. When a virtual disk is created through the New Virtual Hard Disk Wizard, users are given a choice to create a blank disk or copy the contents from a physical disk (as in \\.\PHYSICALDRIVE0). The wizard does not provide any option to copy only a partition within a physical disk nor is there a utility to perform a physical disk copy after the wizard is finished.

VHD files are limited to 2,040 GB in size, 8 GB short of 2 TB. Consider this fact when provisioning a single large disk for a virtual machine. Consider using a passthrough disk that does not have the imposed limitation.

Passthrough disk

Passthrough disk is an alternative to VHD, whereby virtual machines have direct access to the physical disk, similar to Raw Device Mapping (RDM) with VMware. Passthrough is only applicable to block devices such as iSCSI or FC. “[LUN configuration options](#)” on page 23 provides more information about this disk type.

<p>Best practice: In an I/O intensive environment such as an online transaction processing (OLTP)/decision support systems (DSS) database workload, it is strongly recommended that you use fixed-size VHDs or passthrough disks to achieve best performance. Do not use dynamically expanding disks for a production workload because a performance penalty is imposed as a result of allocating data blocks on demand.</p>

Dynamic virtual machine storage

Windows Server 2008 R2 enables VHD files and passthrough disks to be added to and removed from powered-on virtual machines through the virtual SCSI controllers (no IDE controller), without the need to power off the virtual machines. The virtual machines must have Hyper-V R2 Integration Services installed to take advantage of hot-add and hot-remove storage.

Failover cluster storage

In a failover cluster environment, do not use dynamic disks (as opposed to basic disks in disk management) as a shared storage because Microsoft does not support this configuration and it can result in file system corruption.

Basic and dynamic disks

Microsoft recommends the use of basic disks with the SAN storage connected to Windows hosts.

Caution: Do not use dynamic disks in a failover cluster environment because it can result in file system corruption.

For storage subsystems supporting failover cluster in Windows Server 2008, SCSI-3 Persistent Reservation (PR) commands must be supported. Celerra iSCSI PR support is included in release 5.6.39 and later. The *EnableAptpl* server parameter is used to enable/disable the Activate Persist Through Power Loss (APTPL) support for all iSCSI targets on a Data Mover. When APTPL support is enabled, the Data Mover supports requests from a cluster to maintain persistent reservation information on non-volatile storage. By default, APTPL support is disabled.

Best practice: Enable Persistent Reservation on a Celerra Data Mover to support failover clustering. To do this, execute the command `server_param movername -facility iscsi -modify EnableAptpl -value 1`. The Data Mover must be rebooted for the change to take effect.

Hyper-V network concepts

The Hyper-V architecture requires multiple network interfaces configured in the parent partition to support the connectivity requirements of the management console, storage back end, and virtual machines.

The primary network interface usages are:

- **Management console** — Provides remote management services to the parent partition of the server running Hyper-V.
- **IP storage** — Used for all IP storage traffic to establish iSCSI sessions.
- **Virtual machine** — External virtual network/switch provides connections to the guest operating systems.

Best practice: Create a separate gigabit interface on a different subnet for each network usage so as to isolate traffic (for security and performance reasons).

Management console

By default, the Remote Desktop feature is disabled in the parent partition. Ensure that the feature is enabled, and that the Remote Desktop Protocol (RDP) connection is unblocked by the Windows firewall to enable remote administration of the parent partition.

IP storage

Based on the storage protocol, EMC recommends slightly different network topologies and high availability (HA) configurations. When iSCSI is used to store VHD files or create passthrough disks, it is recommended that you take advantage of the Multiple Connection per Session (MCS) feature of the Microsoft iSCSI Software Initiator. “[Celerra iSCSI configuration for Hyper-V](#)” on page 18 provides step-by-step instructions on configuring MCS with Celerra.

MCS provides load-balancing capabilities and eliminates the need to aggregate network ports by using port channeling. It requires two network ports on the server running Hyper-V. You need to assign two separate subnets that match those on the Data Mover side to avoid unnecessary routing. [Figure 5](#) on page 14 shows how two network switches connected through an uplink eliminates the possibility of a single point of failure.

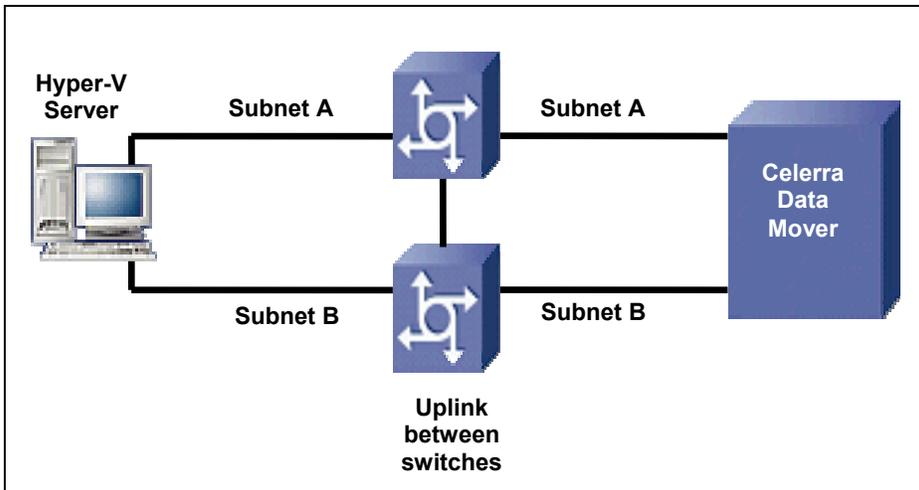


Figure 5 Celerra iSCSI networking with Hyper-V

Best practice: Use one or more gigabit network interfaces to access the IP storage targets.

Virtual machine networks

You can create three types of virtual machine networks with Hyper-V: external, internal, and private. An external network is the only virtual network type that requires binding to a physical network adapter so that virtual machines can access a public network. Virtual networks are configured by using Virtual Network Manager in Hyper-V Manager (Figure 6).

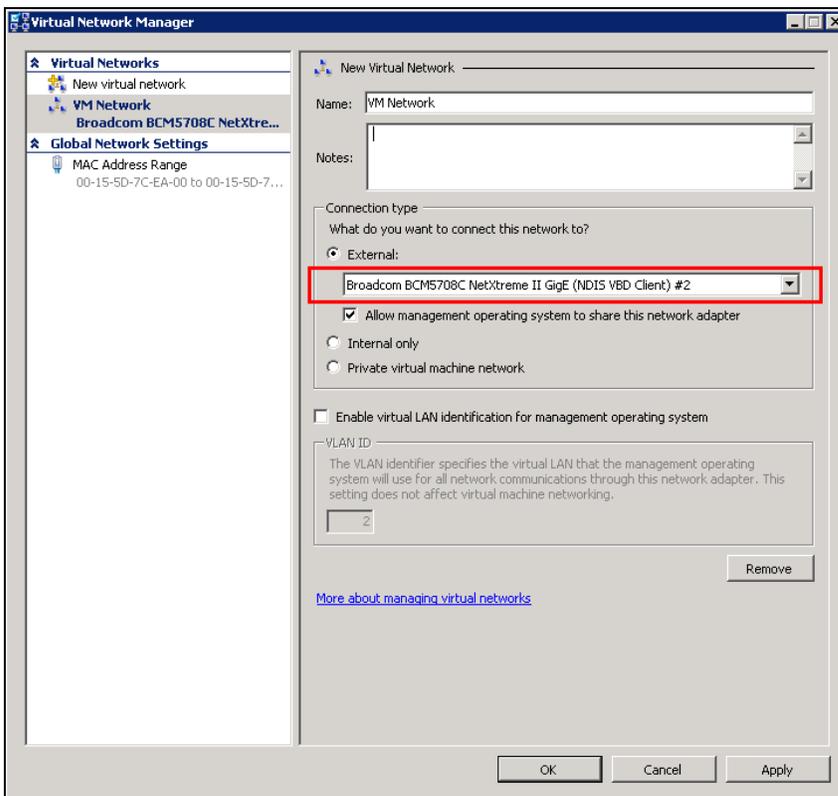


Figure 6 Virtual Network Manager wizard

To configure an external network, select an adapter from a list of available physical interfaces. The physical interface becomes a pseudo adapter that runs the Microsoft Virtual Network Switch Protocol as shown in Figure 7. The interface is listed in the Network Connections window, but it is hidden from the output of the `ipconfig` command because TCP/IP is disabled by default for this type of interface.

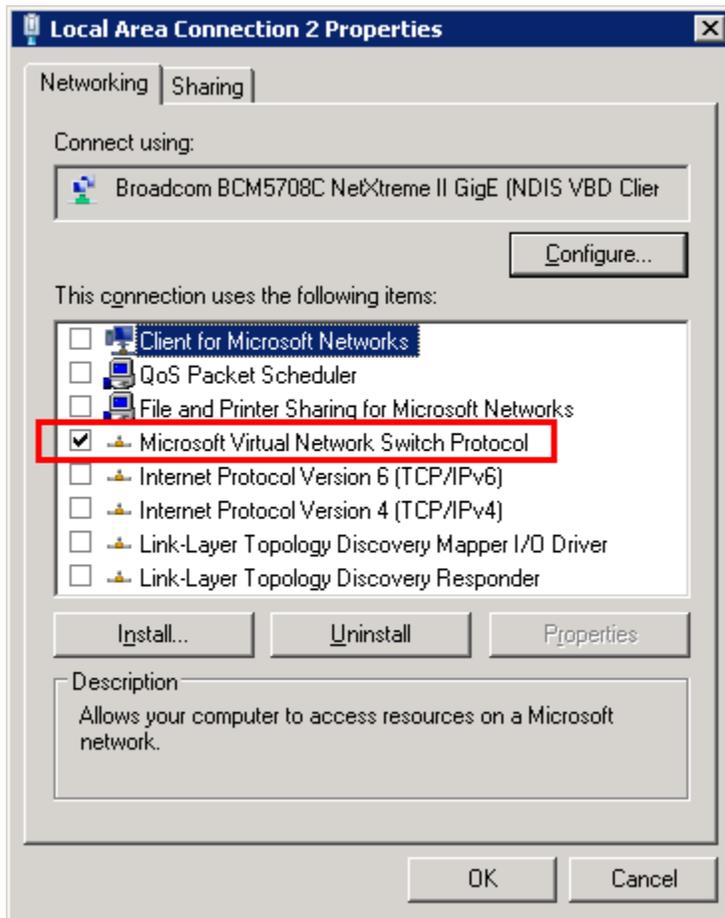


Figure 7 Virtual Network Switch adapter properties

In Figure 7, note that all checkboxes for TCP/IP are cleared. Although you can assign an IP to this interface, it is generally not recommended. Static IP information will be overwritten and completely removed if the virtual network gets recreated.

Best practice: To avoid conflicts, do not share physical network interfaces that are used by external virtual machine networks with the management console or IP storage.

In addition to a pseudo adapter, because the Allow management operating system to share this network adapter checkbox is selected by default (Figure 6 on page 14), Windows creates a new interface (Figure 8 on page 16) through which the virtual machine network traffic travels. Note that the only item that is not selected by default in Figure 8 is Microsoft Virtual Network Switch Protocol. Although you can assign an IP to this interface, it is not recommended for the same reason that is applicable to the pseudo adapter.

If the Allow management operating system to share this network adapter checkbox is not selected, the new interface that the parent partition needs to share will not be created because the physical interface is dedicated to the virtual machine traffic. Only the pseudo adapter is created in this case.

Best practice: It may be difficult to distinguish one interface from another if there are multiple interfaces, especially when new interfaces are added during the virtual network creation. Ensure that you rename and label each interface in **Network Connections** to make it easier to identify their designated purposes.

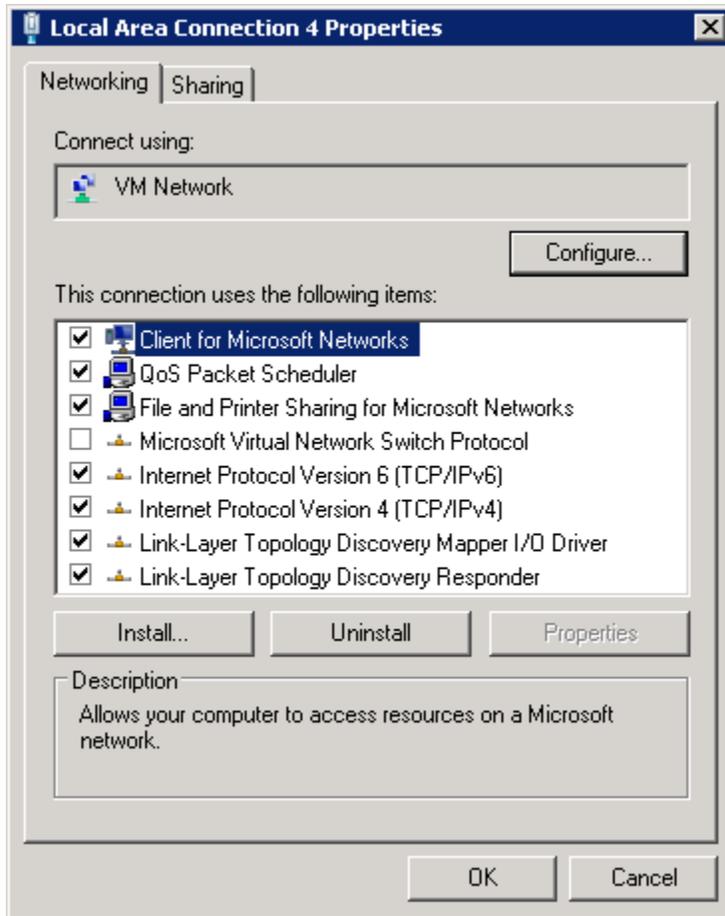


Figure 8 New Interface for virtual machine network traffic

VLANs

It is a recommended practice to segment the network traffic from Celerra through a private LAN by using either a virtual LAN or a dedicated IP SAN switch. Based on the throughput requirements for the virtual machines, you may need to configure more interfaces for additional network paths to the Celerra Data Mover.

VLANs enable secure separation and quality of service (QoS) of the Ethernet packets carrying the Hyper-V server I/O. Their use is a recommended configuration practice for a server running Hyper-V. Hyper-V supports various VLAN options with 802.1Q. You can apply the VLAN tag to the virtual switch network interface to limit the scope of the network traffic on that virtual machine network interface.

The Celerra network stack also provides support for 802.1Q VLAN tagging. A virtual network interface (VLAN) on Celerra is a logical association that you can assign to either a physical or a logical network device such as an aggregated link or a Fail-Safe Network interface. VLAN on Celerra occurs above the Layer 2 hardware address. Therefore, the Data Mover can support multiple physical interfaces for high availability and load balancing.

Celerra and Hyper-V can support an environment in which the storage network has been segmented with VLANs because both provide support for 802.1Q.

Best practice: Segment network traffic for IP storage through a private LAN by using either a VLAN or an independent IP storage area network (SAN) switch.

Jumbo frames

Celerra supports jumbo Ethernet frame sizes up to 9,000 bytes. Based on the NIC capability and its Windows device driver support, enable jumbo frames whenever possible to reduce the host CPU utilization and improve the network efficiency on the network interfaces used for IP storage and on the virtual machine network. Jumbo frames must be supported on all endpoints in the data path including NICs, switches, and target devices. Starting with Windows Server 2008 R2, Hyper-V supports jumbo frames in virtual machines as well as virtual networks. NIC vendor documentation provides further details on jumbo frame support and how to configure it in Windows Server 2008 R2.

Celerra storage provisioning for Hyper-V

With the support for IP storage in servers running Hyper-V, EMC Celerra has become a preferred platform to provision storage used for virtual machines. Celerra iSCSI support provides an option to configure network storage for the servers running Hyper-V. Celerra with CLARiiON® or Symmetrix® back-end arrays provides servers running Hyper-V with highly available, RAID-protected storage.

The use of Celerra as a storage system does not preclude the use of Fibre Channel devices. Servers running Hyper-V support the use of Fibre Channel devices along with iSCSI. When the server has been configured with different classes of shared storage, you can employ virtual machine migration to rebalance virtual disks based on the needs of the application. This offers a method to relocate a virtual machine among the different tiers of storage and balance the needs of the application across the storage systems.

Storage is configured for servers running Hyper-V through the Celerra Management interface. Celerra provides mirrored (RAID 1) and striped (RAID 5) options for the performance and protection of the devices used to create Hyper-V volumes. Celerra Automatic Volume Manager (AVM) runs an internal algorithm that identifies the optimal location of the disks that make up the file system. Storage administrators are only required to select the storage pool type and desired capacity to establish a file system that can be presented to Hyper-V as an iSCSI storage location.

The choice of which storage and RAID algorithm to use is largely based upon the throughput requirements of your applications or virtual machines. RAID 5 provides the most efficient use of disk space with good performance to satisfy the requirements of your applications. RAID 1 provides the best performance at the cost of additional disk capacity to mirror all data in the file system. In testing performed within EMC labs, RAID 5 was chosen for both virtual machine boot disk images and for virtual disk storage used for application data. To identify the appropriate RAID configuration for the servers, it helps to understand the application and storage requirements within the computing environment. EMC NAS specialists are trained to translate those requirements to the correct storage configuration.

The virtual disks are assigned to a virtual machine and are managed by the guest OS just like a standard SCSI device. For example, a virtual disk is assigned to a virtual machine running on a Hyper-V server. To make the device useful, a guest OS is installed on one of the disks. The format of the virtual disk is determined by the guest OS or the install program. A potential configuration is to present an iSCSI LUN from Celerra to a Hyper-V server, which uses the NTFS file system on the iSCSI disk to store the VHD files for a newly defined virtual machine. For a Windows guest, the VHD is formatted as an NTFS file system. You can provision additional virtual disks that are used for applications from one or more Celerra iSCSI LUNs and format them as NTFS by the Windows guest OS.

Naming of storage objects

When you establish a Hyper-V environment with Celerra, you must carefully consider the naming of the storage objects. If you provide descriptive names for the file systems, iSCSI targets, and volume labels in Windows, it can establish valuable information for ongoing administration and troubleshooting of the environment. Prudent use of labels including the storage system from which they are configured is helpful in management and troubleshooting during maintenance.

Best practice: Incorporate identifying elements (for example, iSCSI target names) to the volume label and annotate it with the hostname of Celerra to ease future management and configuration tasks.

Celerra iSCSI configuration for Hyper-V

Unlike Windows 2003 or earlier versions that required the installation of the iSCSI initiator as an add-on component, Windows Server 2008 includes a native iSCSI Software Initiator (accessible from the Control Panel) that you can configure in the Hyper-V parent partition to provision storage for virtual machines. For the configuration of iSCSI, the server running Hyper-V must have a network connection configured for IP storage and the iSCSI initiator service must be enabled. “[Hyper-V network concepts](#)” on page 13 describes the Hyper-V IP storage interface configuration used for iSCSI.

Windows Server 2008 R2 introduces a “Quick Connect” feature for Microsoft iSCSI software initiator, which allows one-click connections to storage devices that do not require advanced settings, such as the use of IPSEC and CHAP authentication. You can use Quick Connect as a one step to perform discovery and login, and to make the target location a favorite location.

To establish a “Quick Connect” iSCSI session to a Celerra in Windows Server 2008 R2:

1. From the **Start** menu, select **Control Panel** and click **iSCSI Initiator**. If the Microsoft iSCSI service is not running, you will be prompted to start the service. Click **Yes** to have the service start automatically each time the computer restarts.

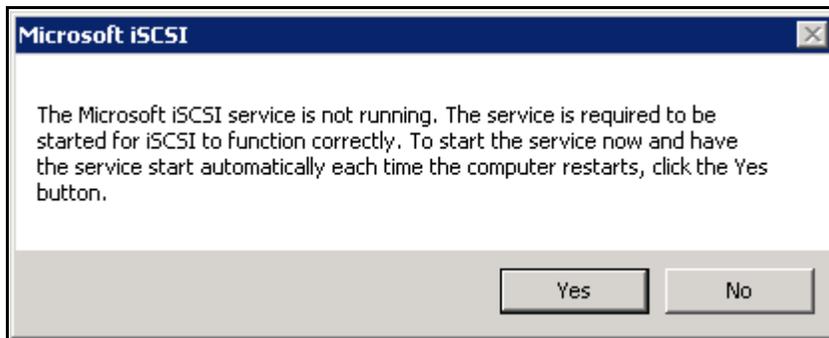


Figure 9 Start Microsoft iSCSI service

2. In the **iSCSI Initiator Properties** dialog box, enter the IP address of the Celerra iSCSI target and click **Quick Connect**.

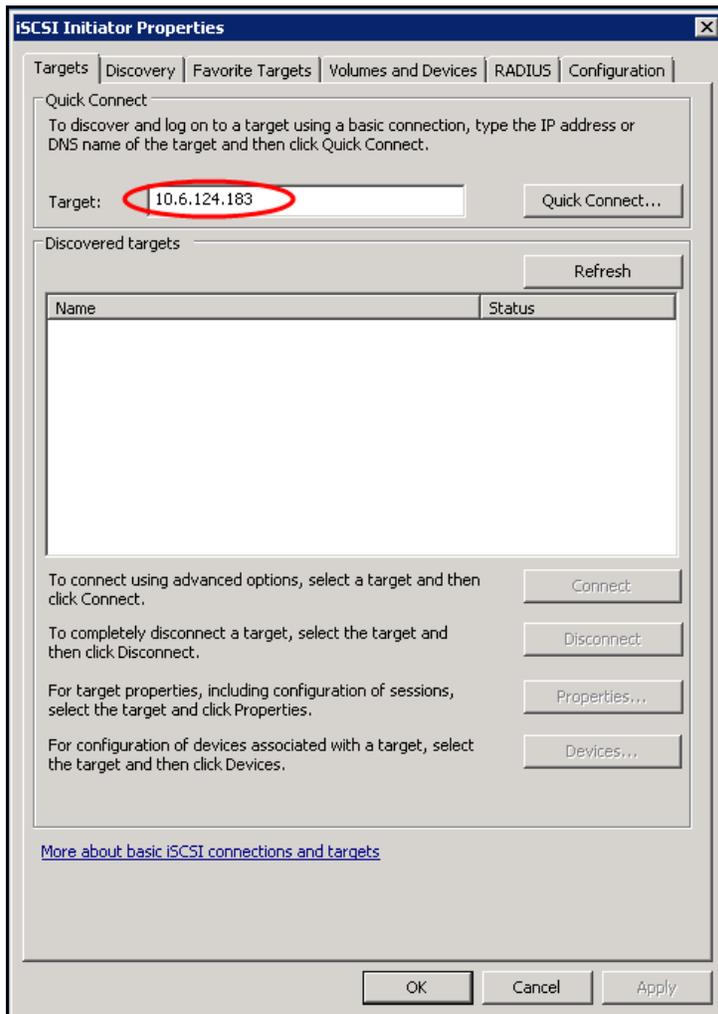


Figure 10 Enter IP address of the iSCSI target

3. If the Celerra has multiple targets as shown in [Figure 11](#) on page 20, select the desired target and click **Connect**. If Celerra has only one target, it is connected automatically.

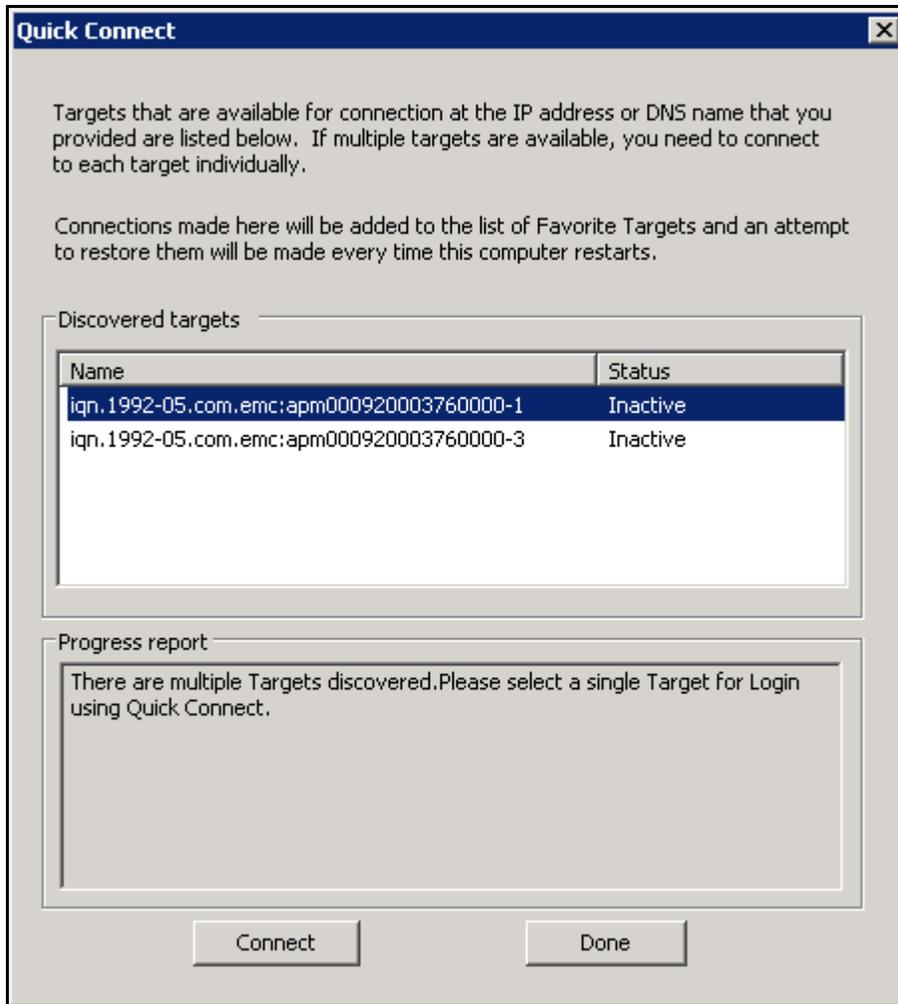


Figure 11 Select the desired target

Before configuring the Hyper-V server for iSCSI, you must ensure that the Windows firewall is modified to enable the iSCSI client in the parent partition. This allows the client to establish sessions with the iSCSI target on Celerra.

Note: By default, iSCSI connections are set as **persistent** across server reboots in Windows Server 2008 R2. Targets are automatically added to the Favorite Targets list so that connections to the Targets list are restored automatically whenever the server reboots.

After an iSCSI connection is established, click **Properties** and confirm that the default connection count is set to one as shown in [Figure 12](#).

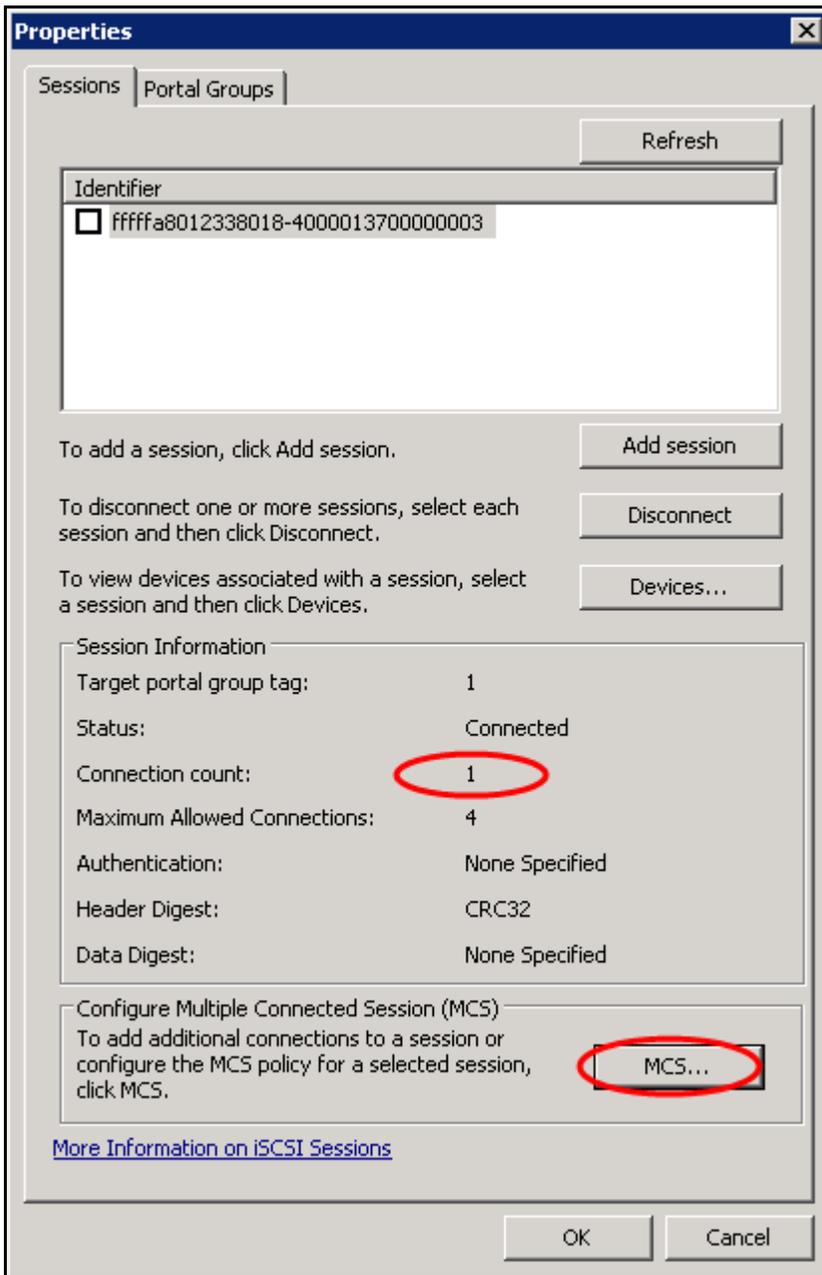


Figure 12 iSCSI Target – Properties

To enable MCS, click **MCS...** and add a second connection as shown in [Figure 13](#) on page 22. Notice that the two connections are established through two separate subnets and the load-balancing policy is set to **Round Robin**.

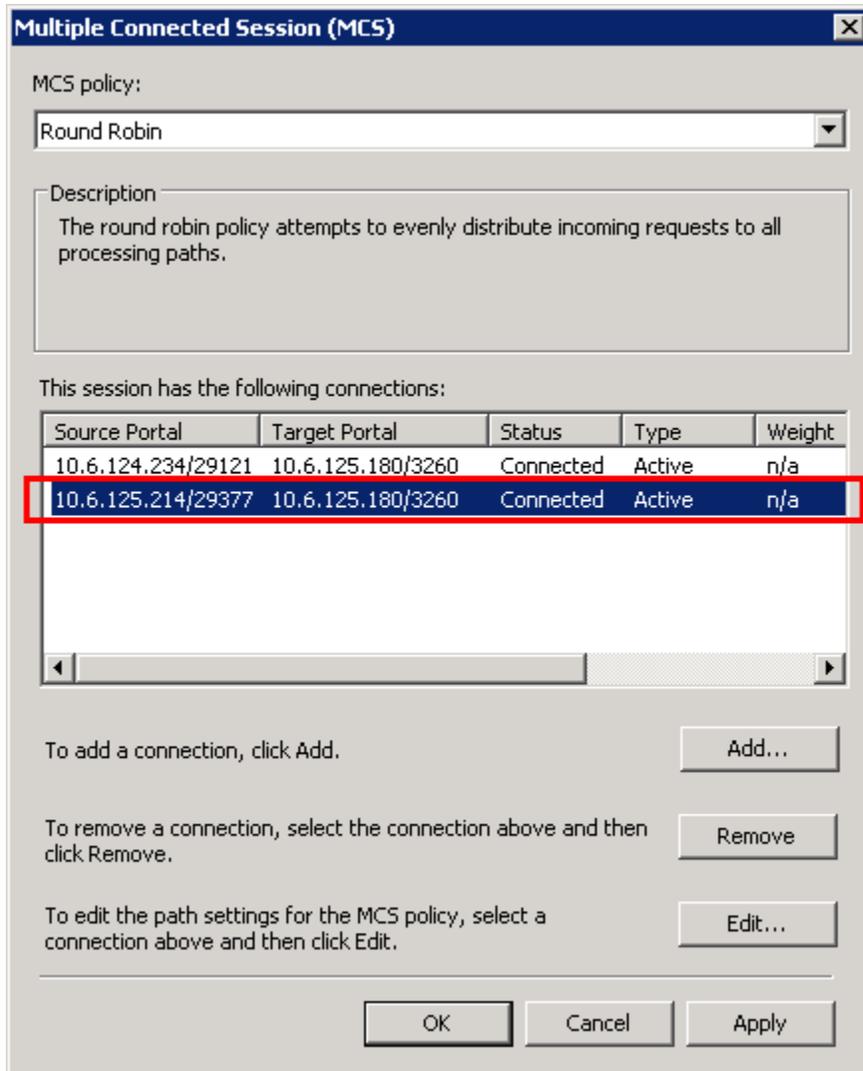


Figure 13 Configuring multiple connections per iSCSI session

Click **OK** to return to the Properties dialog box, and you see that the connection count has incremented to two.

Best practice: For best performance with the EMC Celerra iSCSI target, configure two iSCSI connections per session through two separate subnets and select the Round Robin load-balancing policy.

After you configure the initiator, you need to ensure that Celerra is configured with at least one iSCSI target and LUN.

You can configure an iSCSI LUN in the Celerra Manager iSCSI Wizard as shown in [Figure 14](#) on page 23. If you know the IQN of the Windows host (Hyper-V parent partition) or hardware initiator, you can mask the LUN to the host for further configuration. LUNs are provisioned through Celerra Manager from the Celerra file system and masked to the IQN of the iSCSI software initiator assigned to the Hyper-V parent partition. The IP storage network interface is used to establish the iSCSI session with the Celerra iSCSI target.

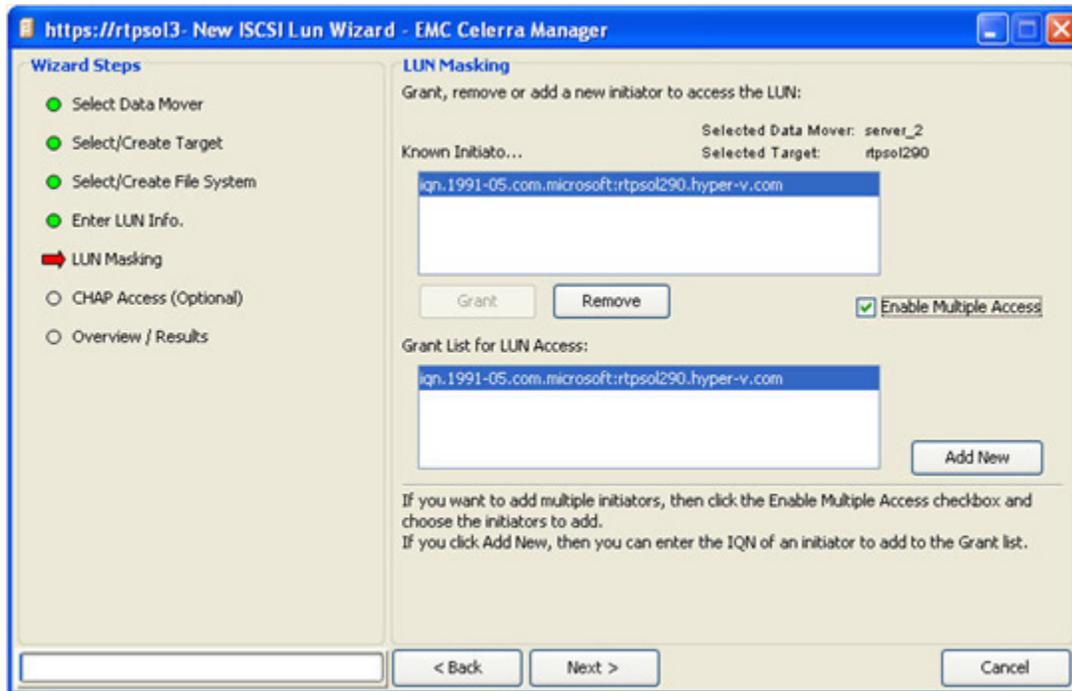


Figure 14 Celerra Manager New iSCSI LUN Wizard - LUN mask assignment

After you complete the configuration steps, open Disk Management in the parent partition and identify the LUNs that have been configured for the Hyper-V server. If the LUNs are not there, you might need to rescan disks to refresh the window.

LUN configuration options

There are three methods to leverage Celerra iSCSI LUNs in a Hyper-V environment: NTFS formatted file system, passthrough disk, and a software initiator (generally Microsoft iSCSI) within the virtual guest OS. Each method is best used in specific use case examples discussed later in this white paper.

VHD

After the Celerra iSCSI LUNs are presented to the server running Hyper-V, you can bring the raw disk online in Disk Management, create a simple volume, format it as NTFS, and assign it a drive letter. Ensure that you use a meaningful volume label to quickly identify the LUNs in the future. You can use the new disk drives to store the virtual machine configuration files, snapshots, and VHD files.

Ideal customer use cases for an NTFS formatted file system are:

- Storage for general virtual machines without I/O-bound workloads such as application servers and Active Directory domain controllers
- Virtual machines that do not require application-level I/O consistency with Native Hyper-V snapshots
- Virtual machines that are ideally backed up by using Visual SourceSafe (VSS)-based backup software

VHD is an excellent choice where various virtual machine LUNs do not require specific I/O performance envelopes or where ease of provisioning and use are paramount.

When Celerra iSCSI LUNs are used with a server running Hyper-V, advanced Celerra features such as virtually provisioned LUNs and dynamic LUN extension are fully supported.

Passthrough disk

Unlike virtual disks that are encapsulated in VHD files on a file system, passthrough disks bypass the parent file system layer and are presented directly to virtual machines as raw physical devices. Because a virtual machine needs exclusive access to the disk, you need to take it offline in the parent partition to avoid unexpected behavior because of interference between the parent and child partitions. Figure 15 shows how the passthrough disk is configured by selecting **Physical hard disk** in the virtual machine settings. You will notice the option is unavailable if the disk in the parent partition is not offline.

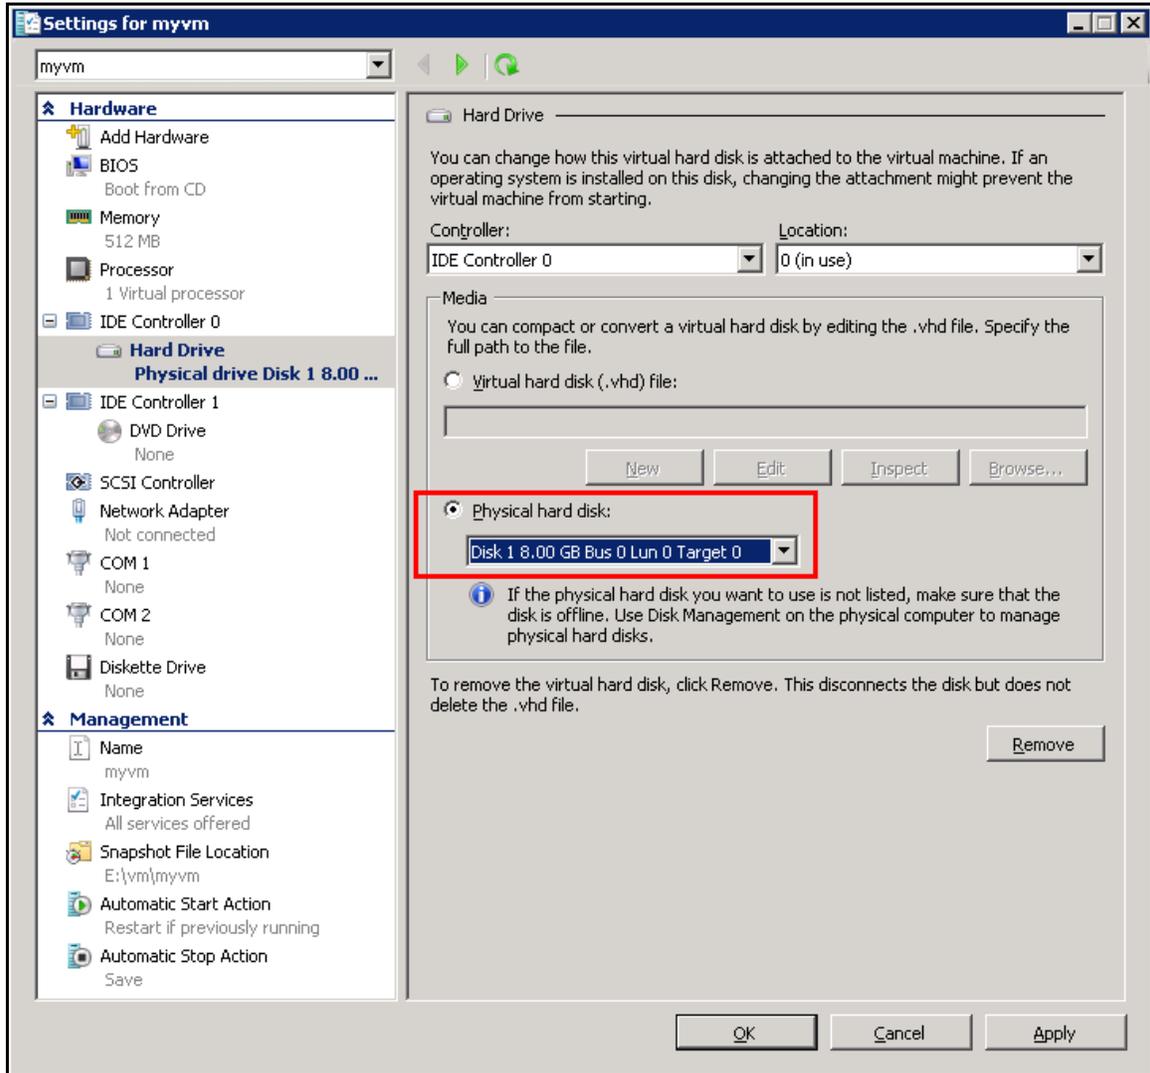


Figure 15 Passthrough disk configuration in the virtual machine settings

The passthrough disk allows you to pass SCSI commands from the guest OS to the SCSI target. The benefit is that you can use management utilities that interface to storage platform-based replication tools to create snaps and replicate iSCSI devices at the array level. However, by default Hyper-V enables SCSI filtering to minimize security risks, so management utilities such as EMC Solutions Enabler (and dependent applications) do not function correctly in virtual machines unless SCSI filtering is disabled. Because of the increased security risk when SCSI filtering is turned off, refer to related EMC and Microsoft Hyper-V documentation to determine whether and how you should proceed to enable storage management for passthrough disks within virtual machines.

Because the passthrough disk does not reside on a file system like VHD, it is not bound by the 2,040 GB size limit as with VHD. On the other hand, the passthrough disk loses some of the benefits of VHDs such as dynamically expanding or differencing VHDs. Another limitation is that, unlike VHD, which can be used to create multiple virtual disks within the file system, you can configure passthrough disks only as a single virtual disk.

Passthrough disks are ideal when a LUN has specific performance needs. Because a passthrough disk has a one-to-one virtual disk to the iSCSI LUN model, it is not shared by multiple virtual machines like VHD. Additionally, application-integrated (for example, Microsoft Exchange 2003 and 2007 VSS) replication is required.

Ideal customer use cases for passthrough disks are:

- Exchange and SQL Server log and database LUNs (you can store boot and binaries in VHD)
- Where application-integrated replication is required

Passthrough disks allow more control of the configuration and layout. However, control and performance tuning implies marginally higher complexity than VHDs. Additionally, you cannot back up virtual machines from the parent partition with passthrough disks by using the VSS-based backup method. You can back up passthrough disks with VSS when initiated from within the virtual machine.

The use of passthrough disks provides the guest OS with the ability to pass SCSI commands directly to the storage device presented by the parent partition. For Celerra, this option is available only if the device is an iSCSI LUN. A motivation for using this device type is when Celerra-level iSCSI replication is required.

The primary difference between a passthrough disk and a file system that hosts VHDs is the number of devices that are supported by them. A file system provides a pool of storage that you can use to create many virtual disks and virtual machines, whereas the passthrough device is presented directly to the virtual machine as a single disk device and cannot be subdivided.

Microsoft iSCSI Software Initiator within a virtual machine

Using the iSCSI initiator within a guest OS is the same as using it on a physical server. The iSCSI initiator is layered on the network stack (in this case the guest OS virtual network), handling all iSCSI tasks in the software as shown in [Figure 16](#).

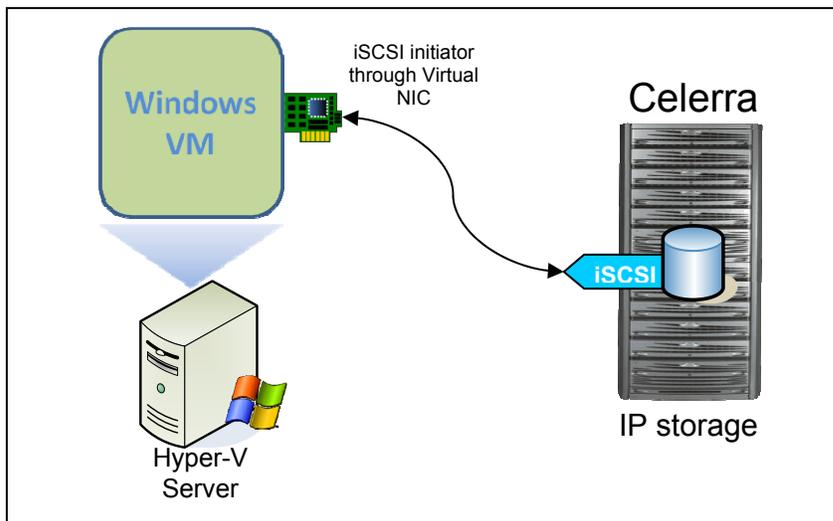


Figure 16 iSCSI initiator access to a Celerra iSCSI target LUN

This use case is similar to the passthrough case because the iSCSI LUN has a one-to-one mapping with the LUN configuration on the physical storage and layout and the performance envelope can be controlled

specifically. However, the primary rationale for this configuration is that this is the only storage model of all three that can support automated handling of the LUN addition and removal within the guest OS. All other methods (NTFS over iSCSI, passthrough) require manual configuration in Hyper-V Manager to configure storage. This capability is critical for mount hosts (for application-consistent replicas) where LUNs are automatically and dynamically added and removed.

The following customer use cases are ideal for using the iSCSI initiator within the virtual machine:

- Mount host for LUN replicas taken using the VSS application interface (for streamed backup)
- Dynamic presentation of SQL Server databases for test and development use

Virtual provisioning

With virtual provisioning (also called thin provisioning) provided through Celerra file systems and iSCSI LUNs, storage is consumed in a more pragmatic manner. Virtual provisioning enables storage devices to be created that do not preallocate backing storage capacity for virtual disk space until the virtual machine application generates some data to the virtual disk. The virtual provisioning model avoids the need to overprovision disks based upon expected growth. Storage devices still represent and support the upper size limits to the host that is accessing them, but in most cases the actual disk usage falls well below the apparent allocated size. The benefit is that, like virtual resources in the Hyper-V server architecture, storage is presented as a set of virtual devices that share from a pool of disk resources. Disk consumption increases based on the needs of the virtual machines in the Hyper-V environment. As a way to address future growth, Celerra monitors the available space and can be configured to automatically extend the backing file system size as the amount of free space decreases.

File system virtual provisioning

Celerra Manager provides the interface to define the virtual-provisioned file system that is used to host the iSCSI LUNs. If you select Virtual Provisioning Enabled when you define a new file system, it creates a file system that consumes only the initial requested capacity (10 GB in [Figure 17](#)). An iSCSI LUN created within the file system realizes the maximum size of 20 GB. As the capacity usage grows, Celerra transparently allocates space up to 20 GB.

New File System	
Create from	<input checked="" type="radio"/> Storage Pool <input type="radio"/> Meta Volume
File System Name:	thin_fs
Storage Pool:	clar_r5_performance 2.2 TB (2358868 MB)
Storage Capacity (MB):	10240
Auto Extend Enabled:	<input checked="" type="checkbox"/>
Virtual Provisioning Enabled:	<input checked="" type="checkbox"/>
High Watermark:	80 % (Optional)
Maximum Capacity (MB):	20480 Required when virtual provisioning is enabled.
Slice Volumes:	<input checked="" type="checkbox"/>
File-level Retention Capability:	<input type="checkbox"/>
Data Mover (R/W):	server_2
Mount Point:	<input checked="" type="radio"/> Default <input type="radio"/> Custom

Figure 17 Celerra virtual-provisioned file system creation interface

Servers running Hyper-V also provide an option to create thin-provisioned virtual disks. By default, virtual disks created using the New Virtual Machine Wizard are thinly provisioned or “dynamically expanding” as Microsoft names it. If you choose not to create “dynamically expanding” disks for performance reasons,

you must select “Attach a virtual hard disk later” when prompted in the wizard, and then add a new virtual hard disk in the virtual machine settings.

Dynamically expanding disks mean that the VHD file space is consumed only when an application or user operating within the guest OS requests a write operation to a virtual disk. The disk request is translated to an allocation request within the Celerra file system. The benefit of this option is that it preserves the amount of storage space used by the virtual disk. However, it does not address the issue of overallocation of the file system. The Celerra virtually provisioned file system improves on this concept by allocating only the disk space it needs to satisfy the requirements of all virtual disks in the file system.

If you view a file system as a one-dimensional expression as illustrated in [Figure 18](#), the non-virtual or allocated file system has a fixed length of 40 GB. In this case, there are several iSCSI LUNs varying in size from 5 GB to 15 GB. If they are thin provisioned, they may not allocate the full complement of their defined space. Therefore, there remains unused space. This space can eventually be used for another LUN or snaps, but presently it remains unused.

In the unallocated storage in [Figure 18](#), a Celerra virtual-provisioned file system is used. In this case, both the server running Hyper-V and the file system preserve disk resources. The file system has a fixed capacity, but it only uses the amount of space that is requested by the virtual machines in the Hyper-V server. The remaining disk storage is free for other file systems. Therefore, the storage that had been allocated in the fixed-length file system is free to be used or shared among the thin-provisioned file systems on Celerra.

Celerra Virtual Provisioning™ further extends the thin-provisioning benefits by allocating only a portion of the disk space and automatically expanding the back up space as the usage requirements increase. This means that Celerra file systems used for Hyper-V consume disk space in a more efficient manner than dense provisioning.

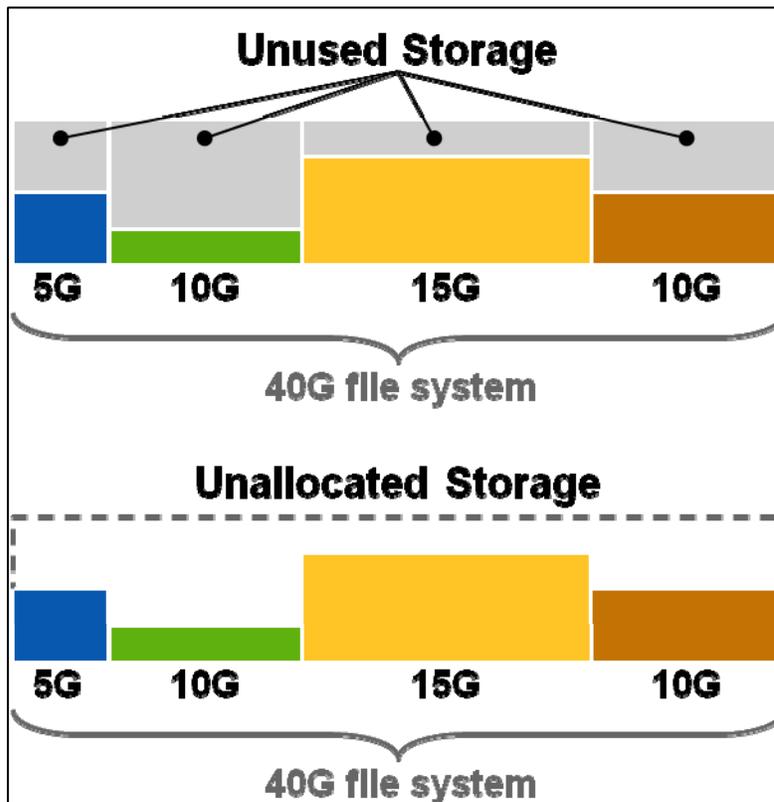


Figure 18 Storage consumption benefits of a Celerra virtual-provisioned file systems

iSCSI virtual provisioning

With Celerra thin-provisioned iSCSI LUN, disk consumption occurs only when the guest OS or application issues a SCSI write request to the LUN. When you create a dynamically expanding virtual disk with Hyper-V in a virtual-provisioned iSCSI LUN, there is little space consumed even after the guest has formatted the virtual disk.

Thin-provisioned LUNs initially reserve no space within the Celerra file system. This Celerra feature provides an effective way to preserve disk resources. As additional virtual machines are added and the disk space becomes full, Celerra uses the auto-extension feature to increase the size of the file system.

When using iSCSI, virtual provisioning provides other options that are tied to the thin-provisioning option on Hyper-V. You can configure the iSCSI LUN as virtually provisioned by specifying the **-vp** option.

When using a virtually provisioned iSCSI LUN, carefully monitor the file system space to ensure that enough space is available. Insufficient space on a file system that contains a virtually provisioned LUN results in data unavailability and may lead to data loss.

Best practice: In addition to monitoring the file system usage, set an event notification for a file system full condition and enable automatic file system extension with a conservative threshold. If enough disk space is not available for the file system extension, data loss can occur.

By using the New Virtual Hard Disk Wizard, you can create a dynamically expanding disk within the thin-provisioned or virtual-provisioned iSCSI LUN. This can be used to create a thin-provisioned NTFS file system as shown in [Figure 19](#) on page 28. The NTFS file system uses the required blocks only when new space requested by the virtual machine is allocated, although Hyper-V interprets the storage as fully allocated. The virtual disk is therefore limited in the amount of blocks consumed on the storage system.

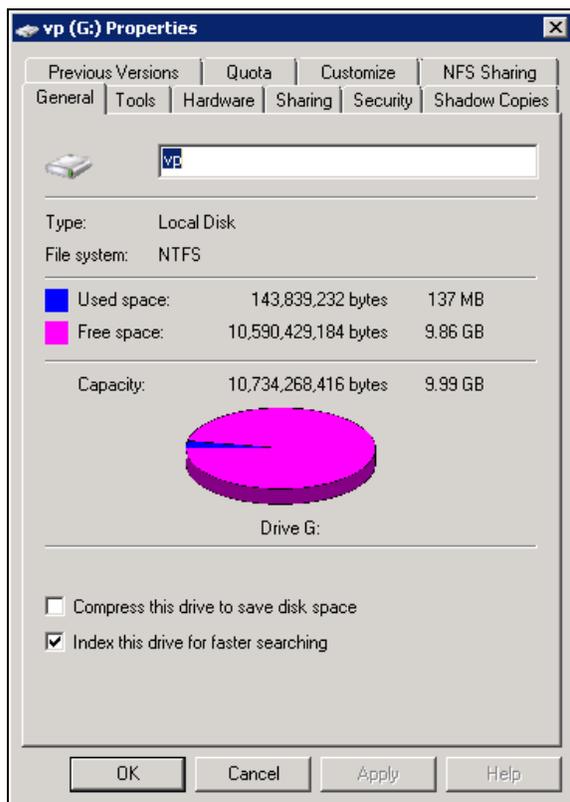


Figure 19 Celerra iSCSI thin-provisioned LUN used as NTFS

When completed and formatted, the amount of consumed space is minimal. [Figure 20](#) shows the File System Properties page on Celerra Manager. This page provides the amount of space consumed within a file system. In this case, the 10 GB LUN has allocated only 146 MB of space after it is formatted as an NTFS file system on a Windows guest OS.

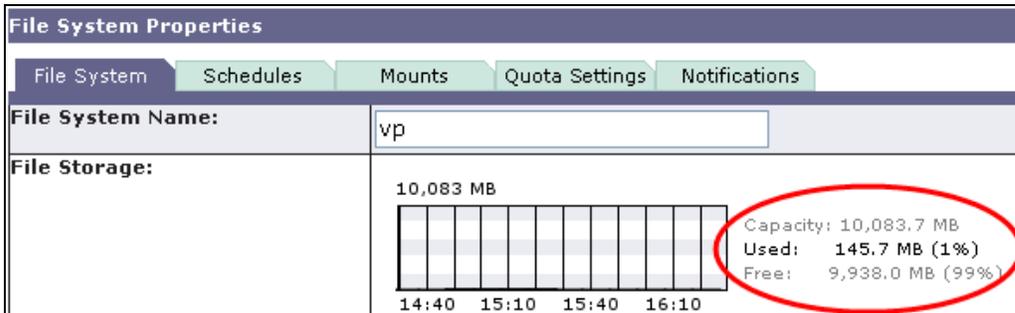


Figure 20 File System Properties page

Partition alignment

Starting with Windows Server 2008, disk partitions are aligned automatically with a 1 MB offset along the proper boundaries by using Disk Management or the diskpart utility. This provides a performance improvement in heavy I/O environments. It is also a good practice to align the disk partitions for the virtual disks within the guest OS. Use fdisk or the partition alignment tools such as diskpart on Windows to ensure that the starting disk sector is aligned to 64 KB boundaries.

iSCSI LUN extension

You can increase the storage space required by a server running Hyper-V by adding additional iSCSI LUNs or extending the size of an existing LUN. The LUN extension is accomplished through Celerra Manager. In [Figure 21](#), an existing 100 GB iSCSI LUN is increased in size to provide additional space to the server. The maximum extension size in this scenario is 1.9 TB. The iSCSI LUN is expanded to 110 GB with a 10 GB extension.

Extend iSCSI LUN	
Data Mover:	server_3
iSCSI Qualified Target Name:	iqn.1992-05.com.emc:apm000746005650000-8
Target:	hyper-v
LUN:	0
File System:	kkfs1
File System Size:	2.7 TB
File System Free Space:	2.5 TB
Current LUN Size:	100.0 GB
Maximum Extension Size:	1.9 TB
Extend LUN By (MB):	10240
LUN Size After Extension:	110 GB
<input type="button" value="OK"/> <input type="button" value="Back"/> <input type="button" value="Apply"/>	

Figure 21 Celerra iSCSI LUN extension window

After a refresh in Disk Management on the Hyper-V server, you see an unallocated space because of the iSCSI LUN extension. If the LUN does not appear in this list, rescan the iSCSI bus.

To adjust the partition boundary and expand the file system, you can right-click the existing NTFS file system and choose **Extend Volume** (this option is available because by default, Windows Server 2008 creates simple volumes) as shown in [Figure 22](#).

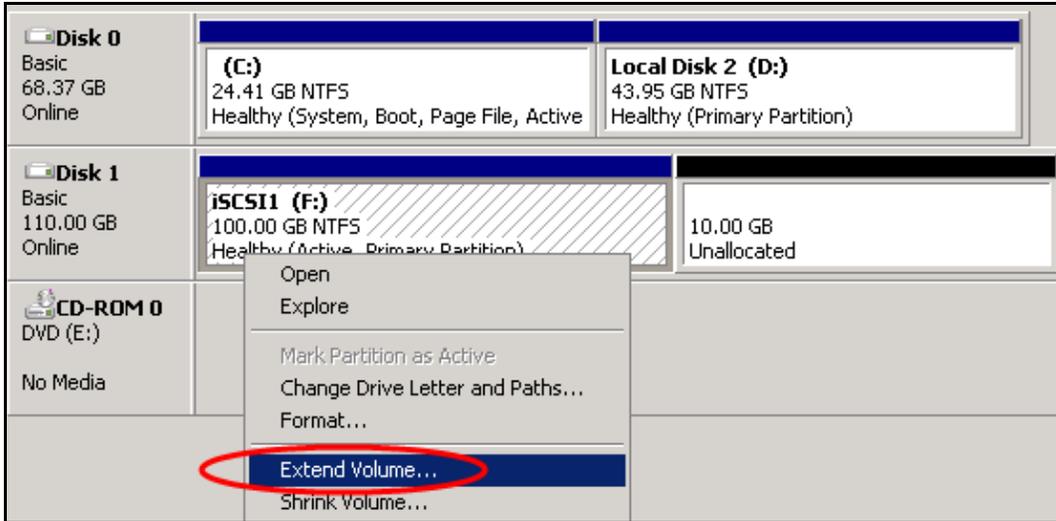


Figure 22 Extend iSCSI disk volume in Disk Management

The **Extend Volume Wizard** in [Figure 23](#) shows the existing device properties in addition to the additional capacity of the iSCSI LUN that can be extended. After the wizard is completed, you see a 110 GB file system as shown in [Figure 24](#) on page 31.

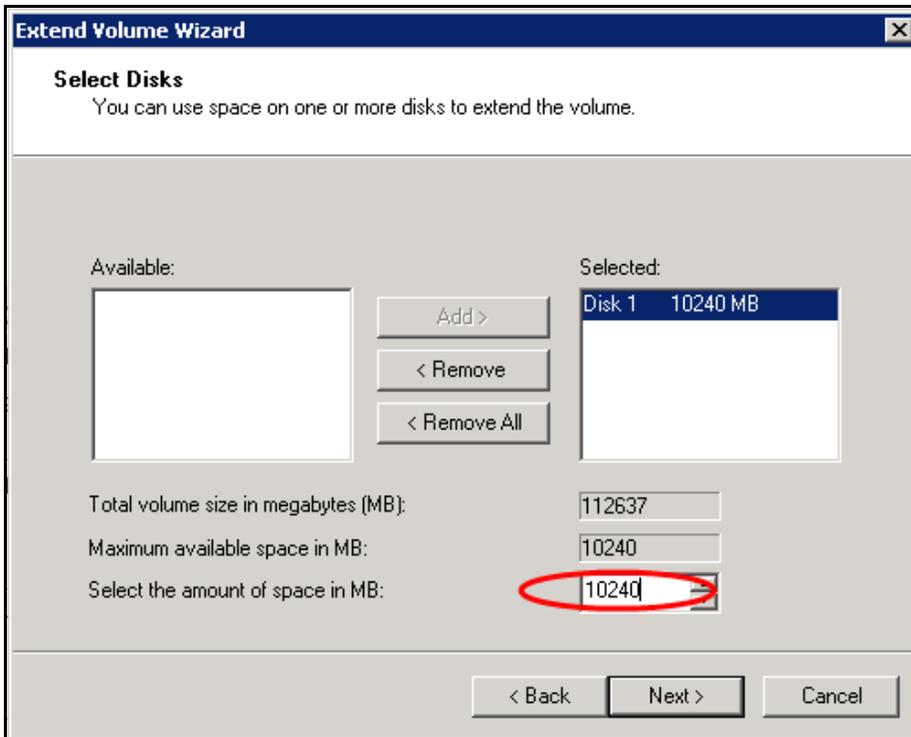


Figure 23 Extend Volume Wizard

Disk 0 Basic 68.37 GB Online	(C:) 24.41 GB NTFS Healthy (System, Boot, Page File, Active)	Local Disk 2 (D:) 43.95 GB NTFS Healthy (Primary Partition)
Disk 1 Basic 110.00 GB Online	iSCSI1 (F:) 110.00 GB NTFS Healthy (Active, Primary Partition)	

Figure 24 Disk Management output after volume extension

Caution: Use caution when extending the iSCSI LUN that is running a guest OS. The LUN extension process may affect running systems as the underlying device is modified. You may need to shut down the guest OS before extending the device in the parent partition.

In a Failover Clustering environment, you can extend the quorum/witness disk, CSV, and non-CSV shared disk by using the same procedure for stand-alone NTFS volumes.

Celerra iSCSI LUN snapshot

Unlike a replicated LUN that requires fully allocated storage space matching the production LUN size, a Celerra iSCSI LUN snapshot is a logical point-in-time copy of the production data that requires only a fraction of the space needed for the production LUN and is created spontaneously. When an iSCSI snapshot is promoted, the snapshot image becomes read/write enabled. It is also known as a temporary writeable snapshot (TWS). Because you can provision TWS rapidly and present it to the same or different host, you can consider the following use cases for Celerra iSCSI snapshot features:

- Test/Dev environment that is created and destroyed spontaneously
- Backup/Restore data stored in a virtual machine
- Snapping iSCSI LUN that is connected through “[Microsoft iSCSI Software Initiator within a virtual machine](#)” as described on page 25.

Test/Dev environment

Although you can take a virtual machine snapshot at the hypervisor level, you cannot power on the virtual machine snapshot alongside the production virtual machine. A workaround is to clone the virtual machine, but it can be time consuming to replicate the associated virtual disk if the size is relatively large.

A promoted iSCSI snapshot or TWS is an ideal solution to this scenario, where a writeable point-in-time copy of the iSCSI LUN hosting the VHDs is created spontaneously and provisioned to a secondary test/dev server as shown in [Figure 25](#) on page 32. When the test/dev environment is no longer required, you can discard the iSCSI snapshot to reclaim the storage space on Celerra.

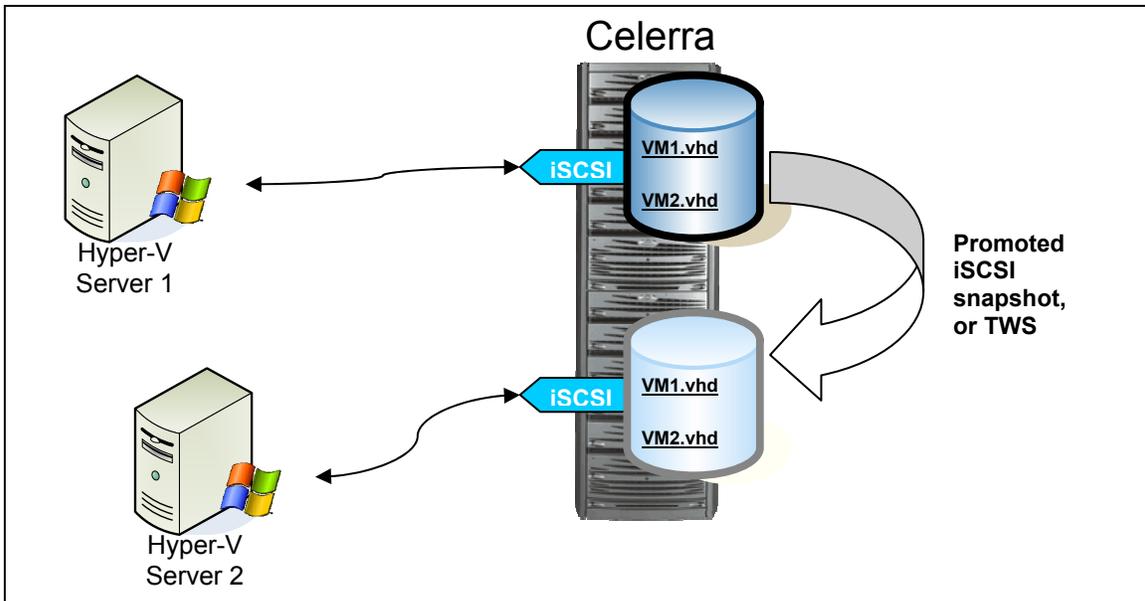


Figure 25 Promoting a Celerra iSCSI snapshot for the secondary server

Best practice: Ensure that there is enough space to store the snapshots because you require extra storage space to create an iSCSI snapshot or TWS. *Sizing Considerations for iSCSI Replication on EMC Celerra —Technical Note on Powerlink®* provides information about how to accurately plan for the required storage capacity.

Backup/Restore data stored in a virtual machine

"Loopback mount" of VHD files was available through the WMI provider in Windows Server 2008 Hyper-V. Windows Server 2008 R2 introduces a new functionality in Disk Management that makes it easy to mount the existing VHD file to view and manipulate its content. In Server Manager, right-click **Disk Management** in Storage, select **Attach VHD** as shown in [Figure 26](#), and specify the location of the VHD file to be mounted in Disk Management.

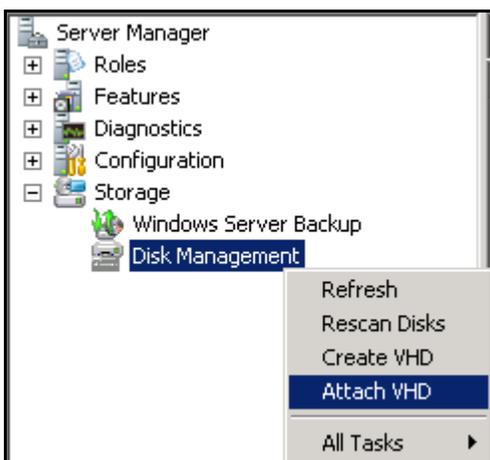


Figure 26 Attach VHD in Disk Management

Figure 27 displays the Disk Management status of a secondary server that is used as a backup server. Disk 4 (E:) resides on a promoted iSCSI snapshot or TWS of a production LUN that is attached to the production host elsewhere. TWS is created and presented to the secondary server because it is likely the production virtual machine is live, so its VHD cannot be mounted when it is in use. You can run a backup agent in the virtual machine but that incurs a performance overhead on the production host during the backup window.

With the TWS attached to the backup server as Disk 4, you can attach or mount the VHD stored in E: in Disk Management as Disk 5. Because Disk 5 is attached as a virtual disk, its icon color is different from the traditional disk icon color. When a drive (F:) is assigned to Disk 5, the backup server can back up the content stored in the virtual machine. You can detach the VHD and destroy the TWS after the backup is completed.

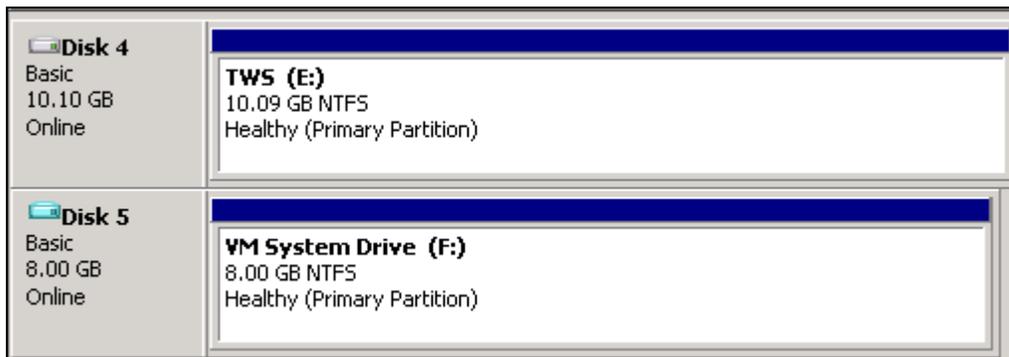


Figure 27 VHD is attached as Disk 5 in Disk Management

Snapping iSCSI LUN that is connected through the Microsoft iSCSI initiator within a virtual machine

Exchange and SQL Server store their databases and logs on Celerra iSCSI LUNs by using Microsoft iSCSI software initiator within a virtual machine, Replication Manager (RM) coordinates with the Microsoft applications to quiesce I/O and leverage the VSS framework to create an application-consistent image by using promoted iSCSI snapshots or TWS. The snapshots are then used in a backup or test/dev solution. The technical validation documents available on Powerlink provide more information on the Microsoft application solutions for Celerra.

Disk signature resolution

In a non-cluster environment, when a production LUN is snapped and promoted as a TWS to the same host, Windows detects two disks with the same signature. As soon as the snapped disk is brought online, Windows resolves the conflict by randomly assigning a new signature to the new disk.

In a cluster environment, if a CSV or a non-CSV shared disk is snapped and promoted to the same host, Windows Disk Management marks the snapped disk as Reserved because it recognizes the same signature as the original clustered disk. Therefore, the snapped disk cannot come online to resolve the signature conflict automatically as compared with a standalone disk.

To work around the duplicated signature issue with a shared disk in a cluster, present the TWS to a host that does not belong to the cluster or that belongs to another cluster. Without prior knowledge of the original disk signature, the new host cannot recognize the snapped disk as a clustered disk. Therefore, it is treated as an ordinary stand-alone disk. After it is brought online, you can use the diskpart command **UNIQUEID DISK** to manually assign a new signature so that it can be presented to the original cluster nodes without any conflict.

Best practice: Avoid presenting TWS to the same host if the production LUN contains a CSV or non-CSV shared disk to prevent a disk signature conflict. Use the diskpart command **UNIQUEID DISK** to display or change the disk signature.

Replication

Celerra provides a framework that you can use with Hyper-V for both local and remote replication as shown in Figure 28. The solution is IP-based and establishes an asynchronous copy of the Hyper-V storage device. When control of the Hyper-V storage device is failed over to the remote Celerra, you can use the copy to restart the Hyper-V environment locally or at a remote location.

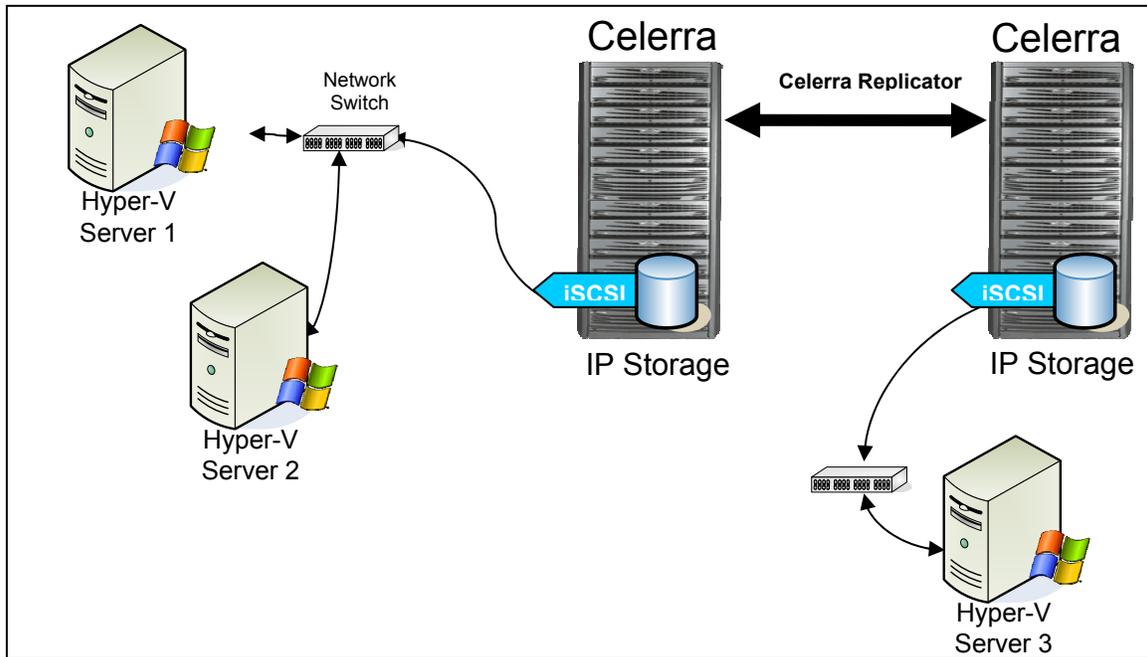


Figure 28 Celerra replication framework

Simple, local, and remote replication of Hyper-V components is supported on Celerra for iSCSI. iSCSI uses block-level replication based on the Celerra IP Replicator feature, which replicates only changed file system or block data for efficient WAN use. Periodic updates between the primary and secondary Celerra provide a crash-consistent iSCSI LUN.

To use the iSCSI LUN on the secondary Celerra, you must set the destination LUN for write access. This is to satisfy the SCSI reservation commands issued to the virtual disks when the guest OS is started. The options to set the file system or LUN to read/write are:

- Fail over the LUN from primary to secondary
- Dissociate the replication relationship between the primary and secondary file systems

Conclusion

To leverage the full capabilities of Hyper-V in Windows Server 2008 R2, you require shared storage for features such as live migration and CSV. When virtualization is used for production applications, the storage performance and total system availability become critical.

The EMC Celerra platform is a high-performance, high-availability IP storage solution that is ideal for supporting Hyper-V server deployments.

This white paper addresses many of the Hyper-V and Celerra considerations. This white paper discusses the use of Celerra iSCSI LUNs to support the creation of local file systems or passthrough disks. Celerra iSCSI LUNs can be:

- Used to create NTFS file systems and CSV in the parent partition
- Accessed exclusively and directly by a virtual machine as raw devices through passthrough disks
- Accessed through the iSCSI software initiator in the guest OS

Aided by the flexibility and advanced functionality of the Celerra products, Hyper-V R2 and Celerra network storage platforms provide a very useful set of tools to establish and support the virtual computing environment.

References

There is a considerable amount of additional information available on the Windows Server 2008 website and discussion boards in addition to numerous other websites focused on hardware virtualization. The Microsoft link <http://www.microsoft.com/windowsserver2008/> also provides additional resources and user guides.