



Compellent Storage Center

Microsoft Hyper-V

Best Practices Guide

Dell Compellent Technical Solutions Group
December 2012

THIS BEST PRACTICES GUIDE IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2012 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the DELL logo, and the DELL badge are trademarks of Dell Inc. Microsoft® and Windows® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims any proprietary interest in the marks and names of others.

Table of Contents

1	Preface	6
1.1	Audience.....	6
1.2	Purpose.....	6
1.3	Customer Support	6
2	Introduction.....	7
2.1	Dell Compellent Storage Center Overview.....	7
2.2	Microsoft Hyper-V Overview	7
2.3	Hyper-V Version and Feature Comparison.....	7
3	Hyper-V Storage.....	10
3.1	Virtual IDE and Virtual SCSI Devices	10
3.2	Virtual Hard Disks	11
3.2.1	Choose a Virtual Hard Disk Type	12
3.2.2	Comparing Fixed and Dynamic VHD/VHDX Files.....	12
3.2.3	Virtual Hard Disks and Thin Provisioning with Dell Compellent	14
3.2.4	Overprovisioning with Dynamic VHD/VHDX Files.....	14
3.3	Options for Presenting Volumes to Hyper-V Guest VMs	16
3.3.1	Pass-through and Directly-attached Disks.....	17
3.3.2	Advantages of Pass-through and Directly-attached Disks	17
3.3.3	Disadvantages of Using Pass-through or Directly-attached Disks	18
3.3.4	Special Considerations When Using Dell Compellent Replay Manager.....	18
3.4	Manage Available LUNs.....	19
3.4.1	Manage LUNs on a Hyper-V Cluster	20
3.4.2	LUN Limits for Large Hyper-V Clusters	20
3.4.3	Avoid the Functional and Physical LUN Limits on Large Hyper Clusters	21
3.5	Single or Multiple Guest VMs per Volume or CSV	22
3.5.1	Advantages of Single Guest VM per Volume	22
3.5.2	Disadvantages of Single Guest VM per Volume or CSV	22
3.5.3	Advantages of Multiple Guest VMs per Volume or CSV	22
3.5.4	Disadvantages of Multiple Guest VMs per Volume or CSV	22
3.6	Mount Points.....	23

3.6.1	CSV Mount Points.....	24
3.6.2	Renaming CSV Mount Points with Server 2012	24
3.7	Disable Automount.....	24
3.8	Cluster Shared Volumes (CSVs).....	25
3.8.1	Advantages of Using Clustered Shared Volumes in a Failover Cluster.....	27
3.8.2	CSV Enhancements with Server 2012.....	27
4	Multipath I/O.....	28
4.1	Install Windows Server Multipath I/O (MPIO)	28
4.1.1	Install MPIO on a Windows 2008 or 2012 Server (with a GUI).....	28
4.1.2	Install MPIO on a Server 2008 Core Installation.....	30
4.1.3	Install MPIO on a 2012 Server Core Installation.....	32
4.2	Configure MPIO Load Balancing Policy.....	36
4.2.1	Change the Load Balancing Policy on a Volume	37
4.2.2	Change the Load Balancing Policy on Volumes by Using MPCLAIM	38
4.2.3	Configuring Server 2012 MPIO Settings with PowerShell.....	39
4.3	iSCSI Software Initiator MPIO Configuration	41
4.3.1	iSCSI Overview	41
4.3.2	Configure a Server 2008 R2/2012 Hyper-V Host for iSCSI MPIO	41
4.3.3	Configure the Dell Compellent Storage Center	46
4.3.4	Configure MPIO for the iSCSI Volume on the Hyper-V Server.....	48
5	Data Instant Replay	52
5.1	Using Replays to Recover Guest VMs	52
5.1.1	Recover a Hyper-V Guest VM with One-to-One Volume Configuration.....	53
5.1.2	Recover a Hyper-V Guest VM with a Many-to-One Volume Mapping.....	57
5.1.3	Recover a Hyper-V Guest VM that resides on a Cluster Shared Volume.....	64
5.1.4	Procedure for Assigning a New Disk ID to a View Volume.....	65
5.2	Use Storage Center Replays to Create a Test Environment	67
5.3	Use Storage Center Replays to Create New Hyper-V Guests from a Gold Image.....	67
5.3.1	Gold Images and Preserving Balanced Controllers.....	76
5.4	Using Gold Images to Deploy Guests to a Hyper-V Cluster.....	79
6	Dynamic Capacity	81
6.1	Dynamic Capacity with Hyper-V.....	81

7	Data Progression with Hyper-V	83
7.1	When to Modify Data Progression Settings with Hyper-V	84
8	Disk Space Recovery with Hyper-V	86
8.1	Server 2012 Hyper-V and Disk Space Recovery.....	86
8.2	Server 2008 R2 Hyper-V and Disk Space Recovery	86
8.3	Configure Disk Space Recovery with Server 2008 R2 Hyper-V.....	87
8.3.1	Disk Space Recovery Example for Server 2008 R2 Hyper-V.....	88
9	Hyper-V Performance Tuning	90
9.1	Storage I/O Performance	90
9.1.1	Block Alignment (Diskpart).....	91
9.1.2	Disable File Last Access Time Check	91
9.1.3	I/O Balancer Controls	91
9.2	Hyper-V Guest Performance	92
9.2.1	Enlightened I/O Guests	92
9.2.2	Integration Services.....	92
9.2.3	Installing or Updating Integration Services.....	94
9.2.4	Background Activity.....	95
9.2.5	Windows Page File.....	96
10	Boot from SAN for Hyper-V.....	97
10.1	Advantages of Booting from Dell Compellent SAN.....	97
10.2	Advantages of Booting from Local Disk.....	97
10.2.1	Domain Controllers on Hyper-V Clusters.....	97
10.3	Configuring Hyper-V Hosts to Boot from SAN.....	98
10.3.1	Configure Boot from SAN with Fiber Channel HBAs	98
10.3.2	Configure a QLogic iSCSI HBA Card to Boot from SAN.....	107
11	Storage Manager for SANs.....	113
11.1	Install Storage Manager for SANs.....	113
11.2	Install Dell Compellent VDS (Virtual Disk Service).....	113
11.2.1	Configure Dell Compellent VDS Manager Service	114
11.3	Storage Manager For SANs Example	115
12	PowerShell Integration	121
12.1	PowerShell Integration Example	121

13	Dynamic Memory	123
13.1	Dynamic Memory Configuration.....	123
13.2	Memory Settings for Hyper-V Guests Running Citrix.....	124
13.3	Memory Settings for Hyper-V Guests running SQL.....	124
14	Virtual Fiber Channel.....	126
14.1	Configure a Hyper-V Guest to use Virtual Fiber Channel.....	127
14.1.1	Create Virtual SANs on the Hyper-V Nodes	128
14.1.2	Prepare a Guest VM to Use Virtual Fiber Channel Adapters	129
14.1.3	Configure Zones for Fiber Fabrics	133
14.1.4	Create a Guest VM Server Object on Dell Compellent	134
14.1.5	Create and Assign a SAN Volume to the Guest VM	138
14.2	Guest VM Clustering with Dell Compellent	142
15	Offline Data Transfer with Server 2012.....	144
16	Live Migration with Server 2012.....	145
17	Hyper-V Replica with Server 2012.....	146
18	Conclusion.....	147
19	Additional Resources	148

Document Revisions

Date	Revision	Author	Comments
06/04/2009	1.0		Initial release
08/11/2009	2.0	Josh Raw	Updated for Server 2008 R2
09/28/2011	3.0	Marty Glaser	Updated for Server 2008 R2 SP1 and Storage Center 5.5
10/29/2012	4.0	Marty Glaser	Updated for Server 2012, Enterprise Manager 6.2 and Storage Center 6.2
12/21/2012	4.1	Marty Glaser	Updated to include additional best practices for virtual fiber channel and guest VM clustering

1 Preface

1.1 Audience

This document is highly technical and is intended for storage and server administrators interested in learning more about how Microsoft Hyper-V integrates with the Dell Compellent Storage Center. Readers should have a good working knowledge of Microsoft Hyper-V and the Dell Compellent Storage Center.

1.2 Purpose

This document provides an overview of Hyper-V and introduces best practice guidelines when integrating Microsoft Server 2008 R2 Hyper-V and Microsoft Server 2012 Hyper-V with the Dell Compellent Storage Center.

Please note that the information contained within this document provides general recommendations only. Configurations may vary based upon individual circumstances, environments, and business needs.

1.3 Customer Support

Dell Compellent provides live support at 1-866-EZSTORE (866.397.8673), 24 hours a day, 7 days a week, 365 days a year. For additional support, email Dell Compellent at support@compellent.com. Dell Compellent responds to emails during normal business hours.

2 Introduction

2.1 Dell Compellent Storage Center Overview

The Dell Compellent Storage Center is an enterprise-class storage area network (SAN) that significantly lowers capital expenditures, reduces storage management and administration time, provides continuous data availability and enables storage virtualization. Storage Center's Fluid Data Architecture manages data dynamically at the block-level, maximizing utilization, automating tiered storage, simplifying replication and speeding data recovery.

2.2 Microsoft Hyper-V Overview

Hyper-V is a layer of software that sits between the physical server's hardware layer and the Hyper-V guest virtual machines (VMs). Hyper-V presents hardware resources in a virtualized manner from the host server to the guest VMs. Hyper-V hosts (also referred to as nodes or virtualization servers) can host multiple Hyper-V guest VMs, which are isolated from each other but share the same underlying hardware resources (e.g. processors, memory, networking, and other I/O devices).

Consolidating many traditional physical servers to virtual servers on a single host server has many advantages: increased agility, better resource utilization, increased power efficiency and reduced operational and maintenance costs. In addition, Hyper-V guest VMs and the associated management tools offer greater flexibility for managing resources, balancing load, provisioning systems, and ensuring quick recovery.

2.3 Hyper-V Version and Feature Comparison

As shown in Table 1 below, Hyper-V was first introduced with the Server 2008 operating system (OS). Server 2012 incorporates the 3rd generation of Hyper-V, which includes many new enhancements and features.

Table 1. Hyper-V Version Comparison

Windows Server Version	Hyper-V Version
Server 2008 Hyper-V	1
Server 2008 R2 Hyper-V	2
Server 2012 Hyper-V	3

An overview of many of the major feature enhancements with Server 2012 Hyper-V is included below in Table 2 as a quick reference.

Table 2. Hyper-V Feature Comparison

Feature or option	Server 2008 R2 Hyper-V	Server 2012 Hyper-V
Logical CPUs per physical host	64	320
Logical CPUs per guest VM	4	64
RAM per physical host	1 TB	4 TB
RAM per guest VM	64 GB	1 TB
Guest VMs per host	384	1024
Guest VMs per cluster	1,000	4,000
Physical hosts (nodes) per cluster	16	64
Virtual Hard Disk Formats	VHD (up to 2 TB)	VHDX (up to 64 TB) (supports VHD also)
Virtual Fiber Channel	---	Supported
Guest VMs on Shared File Storage	---	Supported (SMB 3.0)
Offline Data Transfer (ODX)	---	Supported
Live Migration	1 guest VM at a time within the same cluster only	Multiple guests at the same time. Guests can also be live-migrated to different clusters or hosts, and between stand-alone hosts
Live Storage Migration	---	Supported
Server Message Block (SMB)	Version 2	Version 3 – supports Hyper-V VM files on network file shares
PowerShell Automation Support	Limited	150 built-in PowerShell Cmdlets
Dynamic Memory	Supported, but Hyper- V cannot reclaim unused RAM	Supported, unused memory can be reclaimed allowing for higher guest density
Runtime Memory Configuration	---	Supported – allows dynamic RAM changes to be made while guest VMs are running
Hyper-V Replica	---	Supported – allows for asynchronous replication of guests to another location as part of a DR plan
Guest failover priority	Allows for node-order failover only	Configure failover priority for guest VMs
Disk Space Recovery	Supported on some NTFS volumes only (with Server Agent)	Native support for NTFS and VHDX volumes (no Server Agent Required)

For a complete list of new Hyper-V features included with Server 2012, please refer to the references listed at the end of this document under Additional Resources. Many of the new features and enhancements above will be explored in greater detail in the pages that follow.

Note: Unless otherwise stated, the screen captures, examples and best practices below are applicable to both Server 2008 R2 Hyper-V and Server 2012 Hyper-V. Where significant differences exist, they will be pointed out.

3 Hyper-V Storage

The following sections will explain the different types of storage and storage connectivity options available for Microsoft Hyper-V hosts and guest VMs when integrating with a Dell Compellent Storage Center.

3.1 Virtual IDE and Virtual SCSI Devices

Each Hyper-V guest VM supports a maximum of only two virtual IDE controllers. Each of the two virtual IDE controllers supports a maximum of 2 virtual IDE devices, for a total of 4 virtual IDE devices per guest VM.

Each Hyper-V guest VM also supports up to 4 virtual SCSI controllers with up to 64 devices per controller, for a maximum of 256 virtual SCSI devices per guest VM.

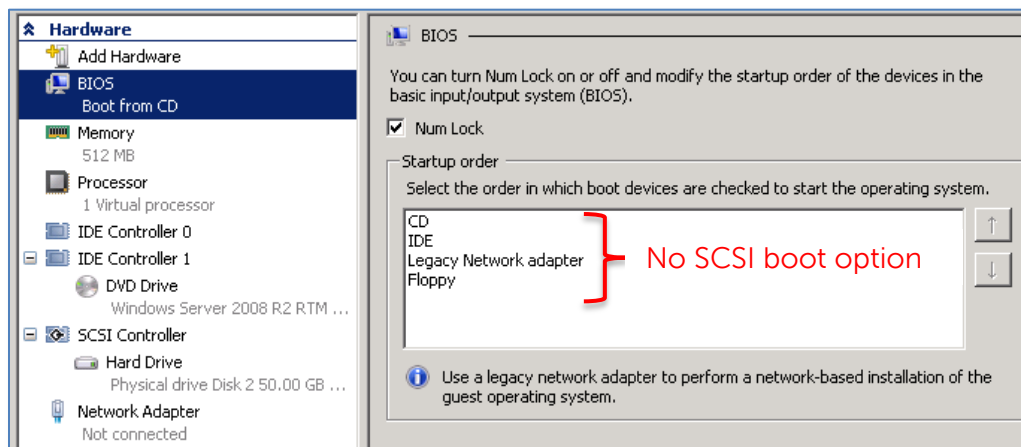


Figure 1. Hyper-V Guest VM Virtual BIOS Startup Options

An important consideration with designing a Hyper-V host and guest VM environment is to understand the boot disk requirements for a Hyper-V guest VM. As shown in Figure 1, when configuring a guest VM’s BIOS settings, the boot volume must be presented to the guest VM as a virtual IDE device (no SCSI option exists). This is a limitation of Microsoft Hyper-V.

Since the guest VM’s boot volume must be an IDE device, this allows for up to three additional virtual IDE devices to be presented to the guest VM. Normally, one virtual IDE device is reserved for use on the guest VM as a virtual CD/DVD drive to mount ISO image files.

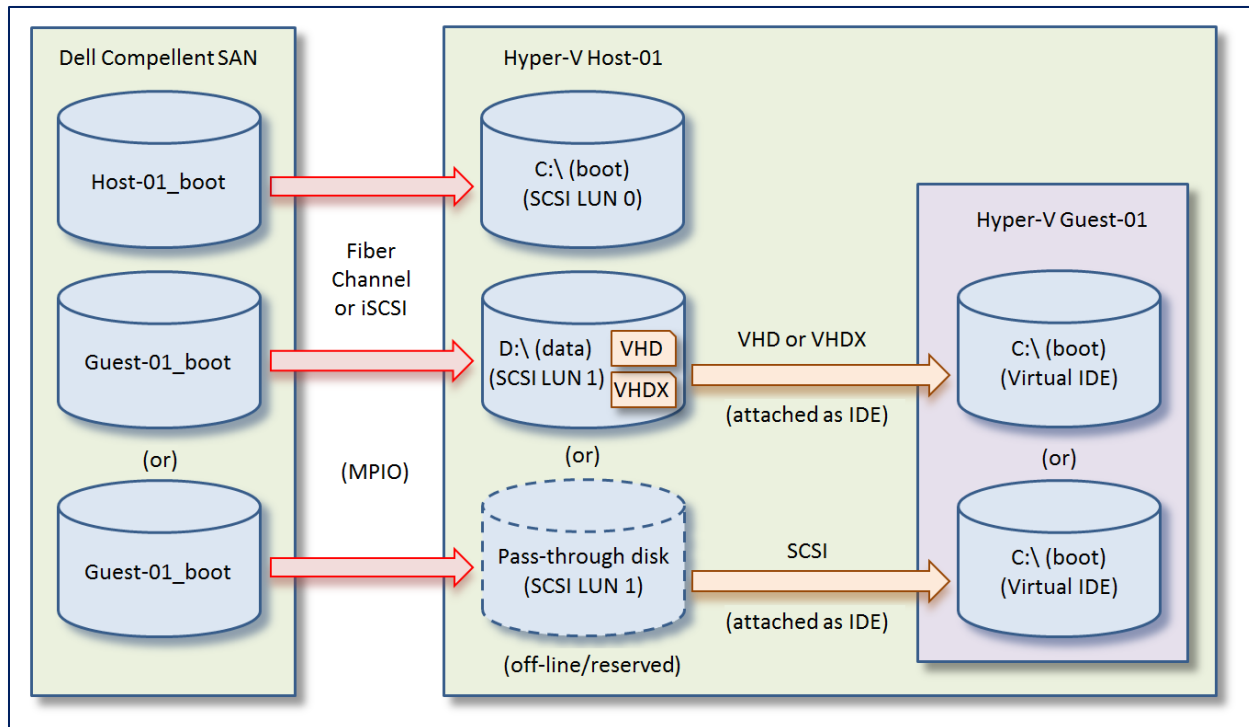


Figure 2. Hyper-V Guest Virtual IDE Boot Device Configuration Options

The IDE boot device that is presented to the guest VM can be:

- A VHD or VHDX hard disk file located on:
 - A SAN volume
 - A local disk on the host server
 - An SMB file share (Server 2012 Hyper-V only).
- A pass-through disk (passed from the SAN through the host to the guest VM).

Note: the VHDX format is supported on Server 2012 Hyper-V only.

Note: Although the I/O performance of physical SCSI and IDE disks can differ significantly, virtualized SCSI and IDE devices in Hyper-V offer nearly identical I/O performance assuming that any required integration services are installed on the guest's operating system. However, it is best practices to use virtual IDE only for the boot drive on Hyper-V guest VMs. For more information on integration services, please refer to section 9.2.2 below.

3.2 Virtual Hard Disks

With Hyper-V, one or more virtual hard disk files can be created on one or more Hyper-V host volumes and presented to a Hyper-V guest VM. A virtual hard disk functions simply as a set of data blocks, stored as a regular Windows file using the host OS's file system. There are two different kinds of virtual hard disk formats:

- VHD is supported with all Hyper-V versions and is limited to a maximum size of 2 TB.
- VHDX is supported with Server 2012 only. VHDX files offer better recoverability in the

event of a power loss, better performance, and support a maximum size of up to 64 TB. VHD files can be converted to VHDX format by using Hyper-V Manager.

3.2.1 Choose a Virtual Hard Disk Type

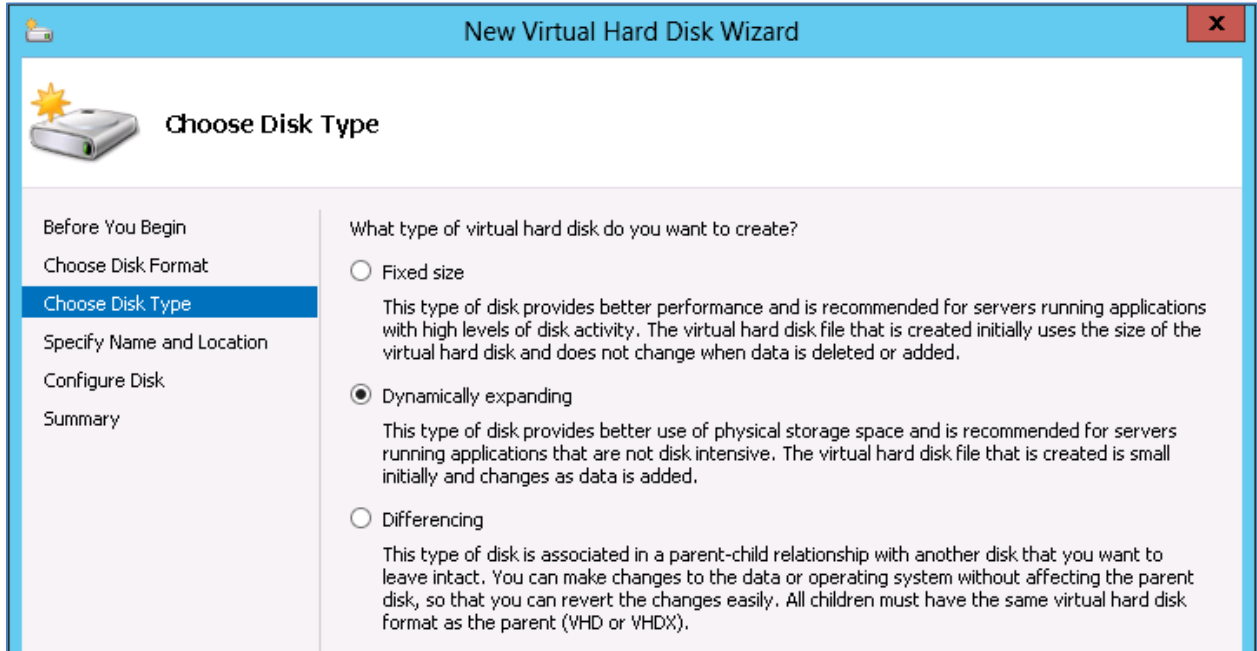


Figure 3. Choose Virtual Disk Type

As shown in Figure 3, there are three different types of virtual hard disk files to choose from when creating either a VHD or VHDX: fixed size, dynamically expanding, and differencing.

3.2.2 Comparing Fixed and Dynamic VHD/VHDX Files

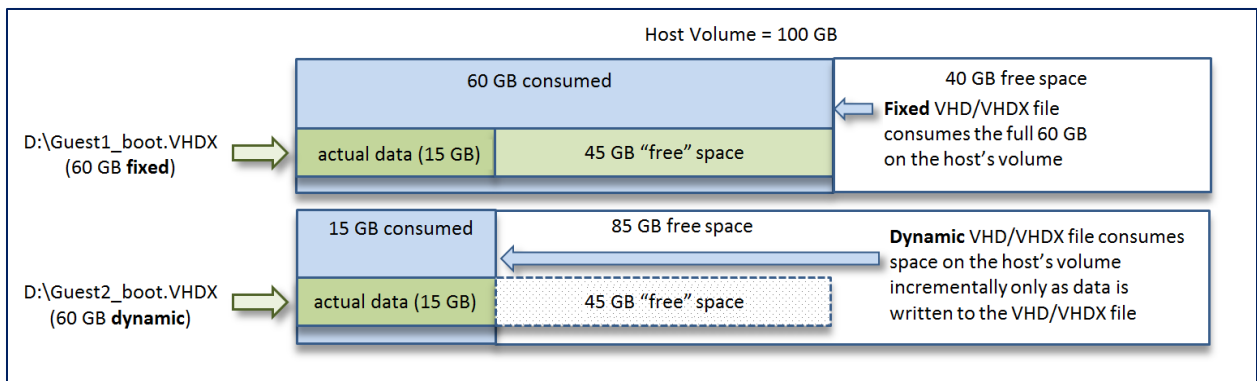


Figure 4. Comparing Fixed and Dynamic Virtual Hard Disks

As shown in Figure 4, a fixed virtual hard disk consumes the full amount of space on the host's volume. A dynamic virtual hard disk only consumes the actual data written to it by the guest VM. From the perspective of the guest VM, either type of virtual hard disk will appear as a full 60 GB volume to the guest.

While Microsoft has made enhancements to the performance of dynamic and differencing VHD files in Windows 2008 R2, and with VHDX files Server 2012, there are some performance and management considerations to keep in mind when choosing the right kind of format for your environment.

- Fixed-size VHD/VHDX files:
 - Recommended for guest VMs with volumes that require high levels of disk activity, such as SQL, Exchange or page files.
 - When formatted, they take up the full amount of space on the host server's volume.
 - Fixed VHD/VHDX files are less susceptible to fragmentation.
 - Copy time for fixed VHD or VHDX files is longer (e.g. from one server to another over the network) because the file size = the formatted size.
 - Large capacity fixed VHD/VHDX files require much more time to provision because the file size = the formatted size.
- Dynamically expanding VHD/VHDX files:
 - Recommended for guest VMs with volumes that require lower levels of disk activity. Dynamically expanding VHDX files require less CPU and I/O overhead as they grow than dynamically expanding VHD files due to improvements made with the VHDX format.
 - Dynamic VHDs/VHDXs are more susceptible to fragmentation.
 - When formatted, they use very little space initially, and expand only as new data is written to them by the guest VMs.
 - Copy time for dynamic VHD or VHDX files is less than fixed (e.g. from one server to another over the network) because the file size = the amount of actual data actually written to the VHD/VHDX file, not the formatted size.
 - Large capacity dynamic VHD/VHDX disks can be provisioned very quickly.
 - Allows for over-provisioning a host's volume (although there are risks – see Figure 5).
- Differencing VHD/VHDX files:
 - Offers storage savings by allowing multiple Hyper-V guest VMs with identical operating systems share a common boot VHD/VHDX file.
 - All children must use the same format (VHD or VHDX) as the parent.
 - Reads of unchanged data reference the differencing virtual hard disk (therefore this data can reside in a lower tier at RAID 5 and take advantage of the cost savings and read performance of RAID 5).
 - Writes (new data) are written to the child virtual hard disk (which by default is written to the highest tier disk at RAID 10 for maximum write performance).
 - Each Hyper-V based snapshot of a Hyper-V guest VM creates a differencing virtual hard disk to freeze changed data since the last snapshot, and new data is then written to a new virtual hard disk file. Creating Hyper-V based snapshots of a Hyper-V guest VM can elevate the CPU usage of storage I/O, but will probably not affect performance noticeably unless the guest VM

- experiences very high I/O demands.
- Maintaining a long chain of Hyper-V based snapshots of a Hyper-V guest VM can negatively affect the performance of the guest because reading from the virtual hard disk can require checking for the requested blocks in many different differencing VHDs/VHDXs. Therefore, it is important to keep Hyper-V based snapshots of Hyper-V guests to a minimum to maintain optimal disk I/O performance.
- Leverage Dell Compellent Storage Center Replays or Replay Manager VSS-aware backups of Hyper-V guests to eliminate the I/O impact caused by taking Hyper-V based snapshots.

Note: With Windows 2008, fixed VHDs will always perform better than dynamic VHDs in most scenarios by roughly 10% to 15% with the exception of 4k writes, where Fixed VHDs will perform significantly better. For Windows Server 2008 R2 and Server 2012, fixed and dynamic VHDs/VHDXs will offer nearly the same performance. However, fixed VHDs/VHDXs are still recommended in cases where disk activity is expected to be extremely high.

3.2.3 Virtual Hard Disks and Thin Provisioning with Dell Compellent

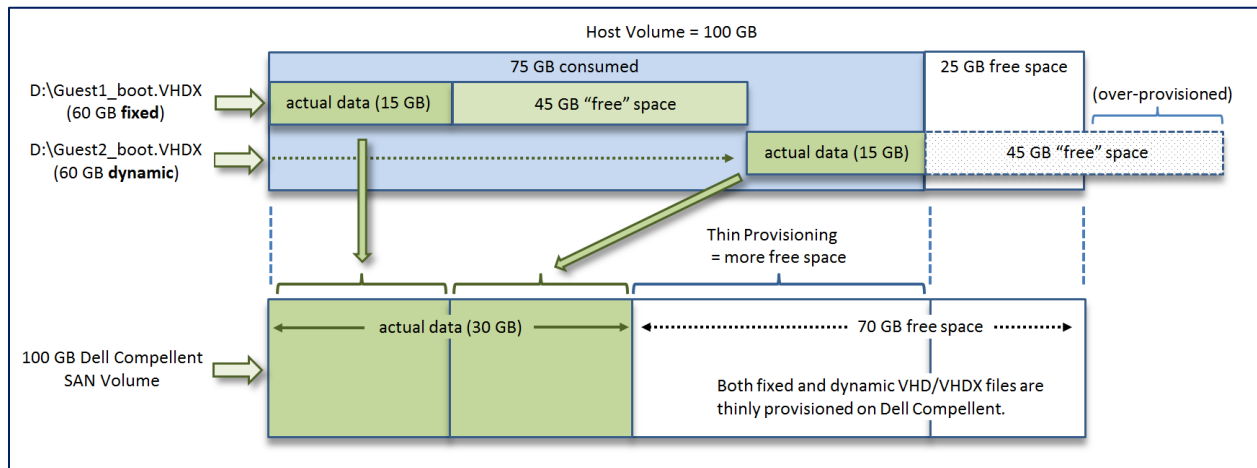


Figure 5. Thin Provisioning of Virtual Hard Disks with Dell Compellent

As shown in Figure 5, it does not matter which type of virtual hard disk format is chosen with regards to thin provisioning and space saving on a Dell Compellent SAN. A 60 GB fixed VHD/VHDX file that has 15 GB of actual data will still only consume 15 GB of space on the Dell Compellent SAN, even though a fixed VHD/VHDX will consume the full amount of space on the volume from the perspective of the host server. Therefore, the performance and management considerations listed above (rather than SAN usage efficiency) would be the main factors to consider when choosing what type of virtual hard disk format to use in your environment.

3.2.4 Overprovisioning with Dynamic VHD/VHDX Files

In Figure 5 above, the example shows how it is possible to overprovision a host's data volume

when using dynamic virtual hard disks. In the example, two virtual hard disk files share the host's 100 GB volume. One virtual hard disk is fixed, and one is dynamic.

While the host's 100 GB volume still has 25 GB of free disk space, guest VM 2 thinks it still has 45GB of "free" disk space. If guest VM 2 added another 25 GB of new data to its virtual hard disk, the host's 100 GB volume would fill up, resulting in a service outage for both guest VMs.

In order to bring the two guest VMs back on line, an administrator would need to expand the host's 100GB volume to allow more room for guest VM 2's dynamic virtual hard disk to expand.

- It is best practices not to overprovision a host's data volume as a space-savings strategy when using dynamic virtual hard disks. It is unnecessary because space savings is already maximized due to Dell Compellent thin provisioning.
- If overprovisioning is used, then ensure that a monitoring system (with alerting) is in place to notify an administrator if the host's volume fills beyond a certain capacity (e.g. 95%). This would allow time for remediation before a service outage occurs.
- If Hyper-V based snapshots are used, allow overhead on the host's volumes for the snapshots.

3.3 Options for Presenting Volumes to Hyper-V Guest VMs

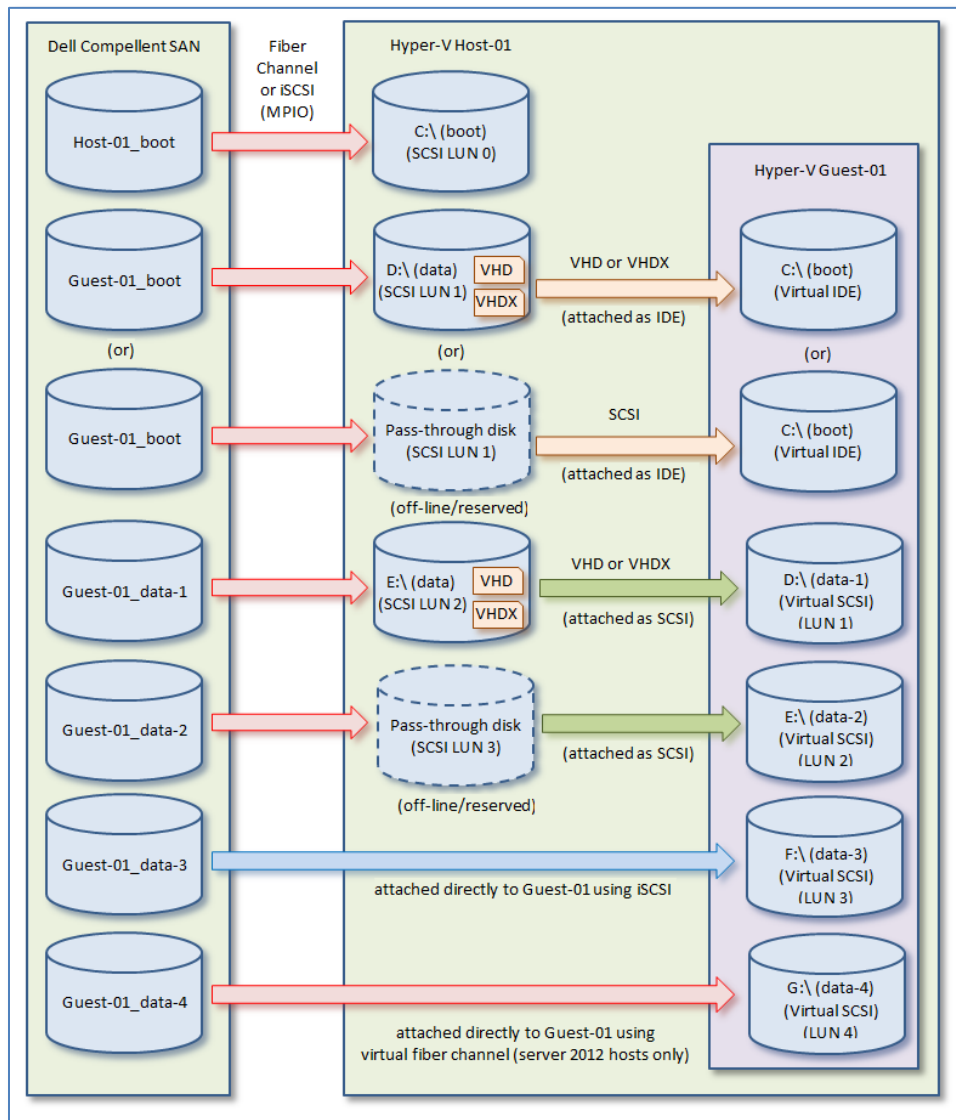


Figure 6. Options for Presenting Data Volumes to Hyper-V Guests

As shown in Figure 6, there are several options available for presenting Dell Compellent SAN volumes to a Hyper-V guest VM. These include virtual hard disks (VHD or VHDX), pass-through disks, iSCSI direct-attached disks, and (new to Server 2012) virtual fiber-channel direct-attached disks. These methods for presenting storage to Hyper-V guest VMs will be discussed below along with some best practice recommendations.

While it is possible to mix and match the different methods in Figure 6 when presenting Dell Compellent SAN volumes to Hyper-V hosts and guests, it is best practices from a management perspective to standardize on a strategy that keeps the environment as simple as possible, and vary from the design only when necessary. For example, the performance gained from a pass-through or directly-attached disk on a guest VM with very high disk activity might be a better choice than a dynamically expanding VHD or VHDX.

3.3.1 Pass-through and Directly-attached Disks

As shown in Figure 6, a pass-through disk is a SAN volume presented to a Hyper-V host that is “passed through” directly to a Hyper-V guest VM. From the host’s perspective the volume is visible, but it is kept in a “reserved” state by Hyper-V.

Directly-attached disks on the other hand, are not visible to (or accessible by) the host server. They are visible to/accessible by the guest VM they are mapped or attached to.

Directly-attached disks are of two types:

- iSCSI: allows a guest VM to attach directly to SAN volumes by configuring the iSCSI initiator software on the guest VM. To present a SAN volume to a Hyper-V guest VM via iSCSI, install the iSCSI initiator software on the guest VM and then configure the iSCSI target (the SAN). Native iSCSI support is provided with Server 2008 and above, and Windows Vista and above.
- Virtual fiber channel: (supported with Server 2012 only) allows physical fiber channel host bus adapters (HBAs) that support the N_Port ID Virtualization (NPIV) protocol on the host server to pass virtual world wide names (WWN)s from virtual fiber channel HBAs on the guest VMs to the fabric. This allows for the creation of guest VM server objects on Dell Compellent to which SAN volumes can be mapped directly.

3.3.2 Advantages of Pass-through and Directly-attached Disks

Pass-through and directly-attached disks offer many of the same advantages:

- Better performance (as compared to a VHD or VHDX) is probably the main consideration for using a pass-through or directly attached disk. They offer better performance because they bypass the host server’s file system, which reduces CPU overhead on the host for managing guest VM I/O).
- With the I/O on a pass-through or directly-attached disk dedicated to a specific guest, it also makes it easier to understand and interpret the I/O demands and usage for that volume using Dell Compellent’s SAN-based I/O reports and tools. By comparison, with multiple VHD/VHDX files on a shared host volume, understanding the I/O has to be done from the perspective of the guest VM’s OS.
- Choosing a pass-through or directly attached disk over VHD or VHDX also avoids the size limitations imposed with VHD or VHDX files.
 - A VHD file (on server 2008 or 2012) is limited to a maximum size of 2040 GB (8 GB short of 2 TB). The VHD size limit is a Microsoft limitation, not a SAN limitation.
 - A VHDX file (supported on 2012 hosts only) supports a much larger size, but is still limited to 64 TB.
 - Pass-through and directly-attached disks can be any size, up to the maximum size that is supported by the guest VM’s operating system.

- Directly-attached iSCSI disks (Server 2008 R2 and Server 2012) and directly-attached virtual fiber channel disks (Server 2012 only) allow Hyper-V guest VMs to be clustered. This can be useful in cases where Hyper-V guest VMs need failover redundancy. This topic is addressed in greater detail in Section 14.2 below.

3.3.3 Disadvantages of Using Pass-through or Directly-attached Disks

There are also a few possible disadvantages to be aware of when using pass-through or directly attached disks:

- The ability to perform native Hyper-V VHD/VHDX snapshots is lost. However, the ability to leverage Storage Center Replays of pass-through and directly-attached storage is unaffected.
- If using a pass-through disk as a boot volume on a guest VM, the ability to use a differencing VHD/VHDX virtual hard disk is lost. However, with the ability to leverage thinly-provisioned Dell Compellent View Volumes based on gold images, the same result (maximizing SAN space utilization) is possible.
- Using pass-through and directly-attached disks may increase the complexity of the environment design resulting in increased management overhead.
- In large environments with a high number of cluster nodes and guests, it is possible to exhaust the pool of available LUN numbers on your hosts when using pass-through disks. This problem is avoided on the host by using directly-attached iSCSI or virtual FC disks. For more information, please refer to Section 3.4 below.

3.3.4 Special Considerations When Using Dell Compellent Replay Manager

- Replay Manager is unable to backup and restore data on pass-through disks. Use Storage Center Replays instead.
- In some cases, it is advantageous to install the Replay Manager 6 agent directly on a Hyper-V guest VM to be able to take more granular backups of SQL or Exchange data. However, these data volumes must be presented to the guest VM as directly-attached iSCSI (Server 2008 R2 and Server 2012) or virtual fiber channel (Server 2012) in order for Replay Manager to be able to backup and restore these volumes.

3.4 Manage Available LUNs

When a Dell Compellent volume is assigned to a Windows Server, the volume is assigned a logical unit number (LUN) by Windows. Windows servers support LUN numbers 0 through 254 (LUN 255 is not supported). Since the Boot LUN in most cases uses LUN 0, that leaves LUNs 1 – 254 available for data volumes.

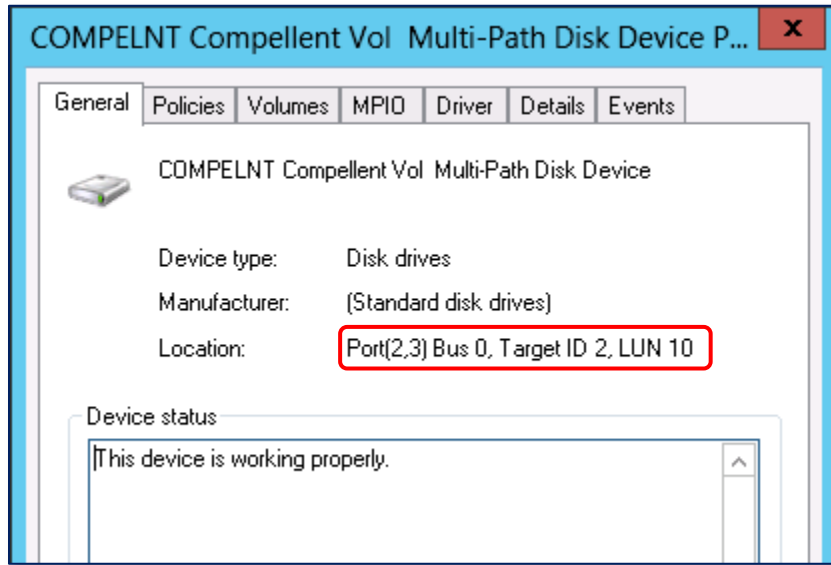


Figure 7. Use Disk Manager to View the LUN, Target, and Bus Details for a Volume

The following table compares the logical LUN limit (as defined by Microsoft) to the total LUNs supported by Dell Compellent, and how the numbers are calculated:

Table 3. Calculating and Comparing LUN Limits

Windows Server "Logical" LUN Limit	Dell Compellent LUN Limitations	
	A Standalone Windows Server	A Windows Server Cluster
254 total LUNs	254 total LUNs	254 total LUNs
128 targets per LUN	2 targets per LUN	1 target per LUN
<u>x 8 busses per HBA</u>	<u>x 1 bus per HBA</u>	<u>x 1 bus per HBA</u>
260,096 total LUNs	508 total LUNs	254 total LUNs

As Table 3 above shows, despite the very high number of LUNs that is theoretically possible, the functional limit of the number of LUNs will be only a fraction of that total (based on performance and capacity variables unique to the hardware and the environment).

As shown in Table 3, Dell Compellent supports up to 508 LUNs for each standalone Windows host server (each LUN number can be used twice), and up to 254 LUNs for each Windows Server Hyper-V cluster (each LUN number can only be used once).

3.4.1 Manage LUNs on a Hyper-V Cluster

In a clustered Hyper-V environment, each cluster shared volume (CSV) must use a unique LUN number that is common to all the nodes in the cluster. For example, if a volume named “CSV_01” is assigned a LUN number of 1 on the 1st node of a Hyper-V cluster, the same LUN number must be used for that CSV when it is mapped to all the other nodes in the cluster.

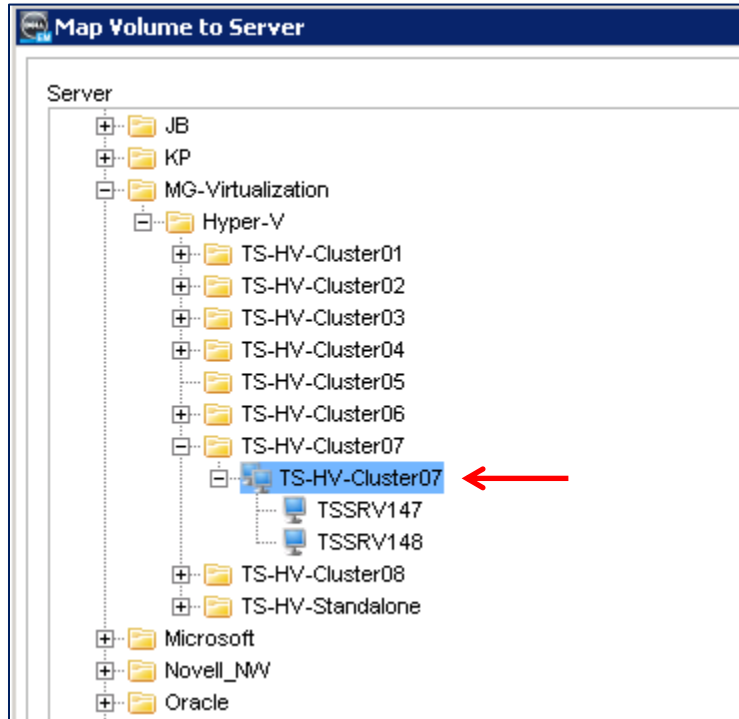


Figure 8. Map Storage to a Cluster Object to Ensure Consistent LUN IDs across All Nodes

As shown in Figure 8, when using Dell Compellent Enterprise Manager or Storage Center Manager to map storage to a Hyper-V cluster object, the process of mapping a new CSV to the next available free LUN number that is common to all the nodes in the cluster is managed automatically.

Once the CSV is mapped, the LUN number for the CSV will be the same on all nodes of the cluster (as viewed from Disk Manager on each node).

3.4.2 LUN Limits for Large Hyper-V Clusters

With large Hyper-V clusters with many nodes, there are two LUN limits to be concerned with:

- The functional LUN Limit: despite the fact that Dell Compellent allows 254 LUNs per Hyper-V cluster, the functional limit of the Servers and HBAs may be reached (their resources may be exhausted) well before the physical limit of 254 LUNs reached.
- The physical LUN Limit: Depending on the Hyper-V cluster design, it is possible to consume many LUN numbers quickly, exhausting the pool of free LUN numbers.

It is also important to note that a small number of available free LUNs numbers must be kept in reserve for an administrator to use for scratch or temporary volumes, and also for SAN maintenance. SAN maintenance might include operations such as expiring Replays or running Replay Manager restore jobs which consume available LUNs temporarily to present View Volumes to the host.

Table 4. Allocation of LUNs on a Hyper-V Cluster

Description	Value
Number of Hyper-V server nodes in the cluster (physical servers configured to boot from SAN)	10
Number of Hyper-V guest VMs on each node in the cluster	12
Total number Hyper-V guests VMs on the cluster (10 nodes x 12 guests on each node)	120
Number of pass-through disks presented to each Hyper-V guest (each disk consumes a LUN number)	2
LUN numbers consumed (120 guest VMs x 2 pass-through disks per guest)	240
LUNs kept in reserve pool (unassigned) for SAN operations (such as Replay Manager Restores or other maintenance)	14
LUN required for each node as a boot volume (LUN 0 - boot from SAN)	1
Total LUNs used on this Cluster (240 + 14 + 1)	255

In the example shown in Table 4, a 10-node Hyper-V cluster with 120 guests VMs using pass-through disks has consumed all but 14 of the pool of available 255 LUN numbers. If additional volumes are needed on these guest VMs, an administrator could present them as directly-attached disks using iSCSI (Server 2008R2 and 2012) or virtual fiber channel (Server 2012 only). Directly-attached disks consume LUN numbers on the guest VMs only - not on the Hyper-V nodes themselves.

3.4.3 Avoid the Functional and Physical LUN Limits on Large Hyper Clusters

To avoid reaching the functional or physical limit of LUNs for a Hyper-V cluster, consider some of the following strategies:

- Use a "many-to-1" guests-to-CSV strategy (using VHD/VHDX files)
- Use direct-attached storage (iSCSI or virtual fiber channel) to present SAN volumes directly to guest VMs.
- Add additional physical fiber channel HBA or iSCSI cards to the cluster nodes (in pairs to ensure MPIO) to increase the number of LUNs available to the nodes (assuming the functional LUN limit has not yet been reached by server OS)
- Create smaller Hyper-V clusters with fewer nodes and fewer guest VMs instead of larger clusters with more guests.

3.5 Single or Multiple Guest VMs per Volume or CSV

One of the design considerations with virtual disk storage is how many Hyper-V guest VMs will exist on a data volume (for a standalone Hyper-V host) or CSV (for clustered Hyper-V nodes). Options include:

- One-to-one configuration (each volume or CSV will host a single guest VM).
- Many-to-one configuration (each volume or CSV will host two or more guest VMs).
- Mix of one-to-one and many-to-one.

There are some design advantages and disadvantages to each strategy depending on the environment.

3.5.1 Advantages of Single Guest VM per Volume

- Easier to track disk I/O patterns per Hyper-V guest VM using the Dell Compellent Storage Center Manager or Enterprise Manager GUI.
- Ability to quickly restore a guest VM by simply replacing the original volume or CSV with a View Volume from a Storage Center Replay.
- If SAN volumes are being replicated to a 2nd location, an administrator has very granular control over which guest VMs get replicated as replication is controlled at the volume (not the guest) level.
- A View Volume (of a Storage Center Replay) can be used as a gold image template to deploy multiple copies of the guest VM to conserve SAN space, or to spin up copies of the guest in an isolated environment for testing or development.
- In server 2008 R2 Hyper-V environments, it is often quicker and easier to manually move a guest VM from one host or cluster to another by moving the volume or CSV rather than copying large VHD files over the network.

3.5.2 Disadvantages of Single Guest VM per Volume or CSV

- More SAN volumes to create and administer.
- Scalability limitations due to the physical and functional limits for the maximum number of LUNs supported on a Hyper-V host or cluster as shown in Table 3.
- Need to create a new volume every time a new guest VM is provisioned.

3.5.3 Advantages of Multiple Guest VMs per Volume or CSV

- Fewer SAN volumes or CSVs to create and administer.
- Easier to provision additional guest VMs due to not having to create new volumes or CSVs each time.
- Avoid the physical and functional LUN number limits for Hyper-V hosts and clusters.

3.5.4 Disadvantages of Multiple Guest VMs per Volume or CSV

- Since replication is done at the volume or CSV level, even if only one guest VM on the

volume or CSV requires replication to a 2nd location, they all get replicated.

- To work around this, designate a separate volume or CSV as a replication volume and only put guest VMs on it that need to be replicated.
- More difficult to manually recover a single guest VM from Storage Center Replay because of having to manually copy one or more VHD/VHDX files back to original location(s). However, The Hyper-V Extension provided with Replay Manager 6 can be used to backup and restore individual Hyper-V guests on a CSV.
- Each time a Replay Manager 6 backup is taken of any Hyper-V guest on a volume, the restore point creates a Replay of the entire volume. This will consume additional SAN space due to extra Replays being created of the entire volume any time any individual guest VM on that volume is backed up.
 - Configure a Replay Manager backup job to back up all the guests on the Volume or CSV as a group to reduce the number of Replays being created of the volume or CSV they share.
- Need to rely on host-based monitoring of the guest VM to get I/O statistics for a single VM.
- SCVMM SAN transfer is not available.
- Live Migration will not work when multiple guest VMs exist on a single volume unless the volume is configured as a CSV.

3.6 Mount Points

Volume mount points are specialized NTFS file system objects which are used to mount and provide an entry point to other volumes without consuming a drive letter. Mount points are created in an empty directory on an NTFS volume.

Mount points are advantageous on Hyper-V hosts because they allow administrators to avoid the 26 drive letter limitation (A – Z) which is an inherent limitation with Windows Server operating systems.

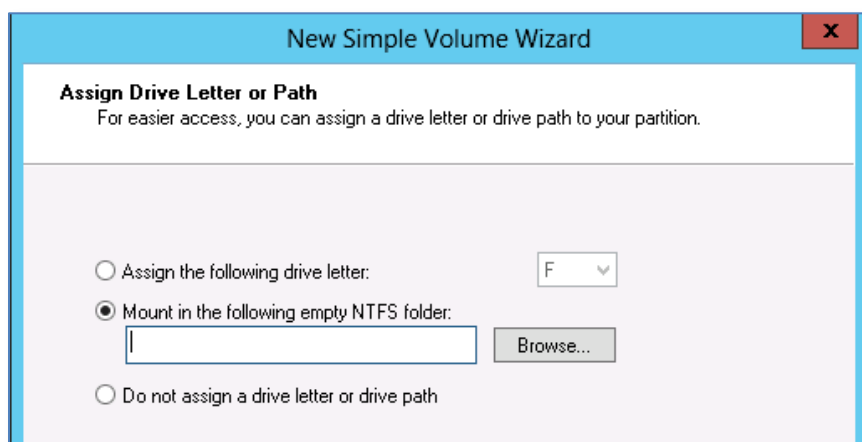


Figure 9. Mount Points Must Use an Empty NTFS folder

The mounted volume itself is not limited to the NTFS file system. Mounted volumes can be formatted with any file system supported by Microsoft Windows, including the new ReFS file system provided with Server 2012.

3.6.1 CSV Mount Points

Cluster shared volumes (CSVs) are presented as mount points by default and do not consume drive letters as they are in an offline/reserved state on the nodes (they are managed by Failover Cluster Manager).

When CSV mount points are created, they are named sequentially as Volume1, Volume2, etc. under the default path (C:\ClusterStorage\) on each node.

3.6.2 Renaming CSV Mount Points with Server 2012

With Server 2008 R2 Hyper-V, CSV mount points can't be renamed. However, with Server 2012, the default CSV mount point names can now be renamed as shown in Figure 10.

Renaming CSV mount points may be desirable for ease of management reasons.

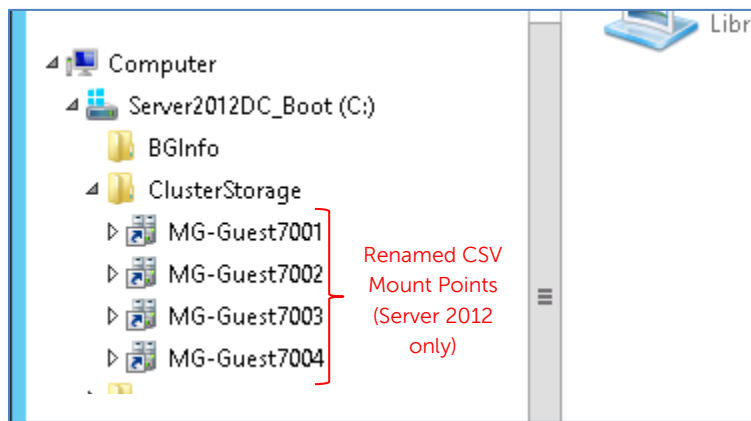


Figure 10. Server 2012 Hyper-V CSV Mount Points Can Be Renamed

Note: rename CSV mount points (Server 2012 Hyper-V) *before* creating guests on them.

3.7 Disable Automount

To prevent the Windows server OS from automatically assigning drive letters to newly mapped volumes, disabling the automount feature is recommended. Having automount disabled is also beneficial when recovering a volume using a Storage Center Replay.

```
C:\Windows\system32\diskpart.exe

Microsoft DiskPart version 6.2.9200
Copyright (C) 1999-2012 Microsoft Corporation.
On computer: ISSRU203
DISKPART> automount
Automatic mounting of new volumes enabled.
DISKPART> automount disable
Automatic mounting of new volumes disabled.
DISKPART> _
```

Figure 11: Disable Automount

To disable automount:

- 1) Open a command prompt window.
- 2) Type **Diskpart** and press **Enter**.
- 3) Type **automount** and press **Enter** to verify the current state of automount. In this example, automount is enabled.
- 4) Type **automount disable** and press **Enter**.
- 5) Verify that the command completes successfully.
- 6) Type **exit** and press **Enter** to close out of Diskpart.

3.8 Cluster Shared Volumes (CSVs)

Cluster Shared Volumes (CSVs) were introduced with Server 2008 R2 and continue to be a core component of Hyper-V clustering with Server 2012. CSVs provide many benefits in a Hyper-V environment.

- CSVs allow many different Hyper-V guest VMs on many different cluster nodes to all read and write to their respective VHD/VHDX files at the same time, regardless of which node currently “owns” the CSV volume.
- CSVs also allow Hyper-V guest VMs to live-migrate or fail over independently to any other node in the cluster without impacting any other guest VMs that might be sharing that same CSV volume.

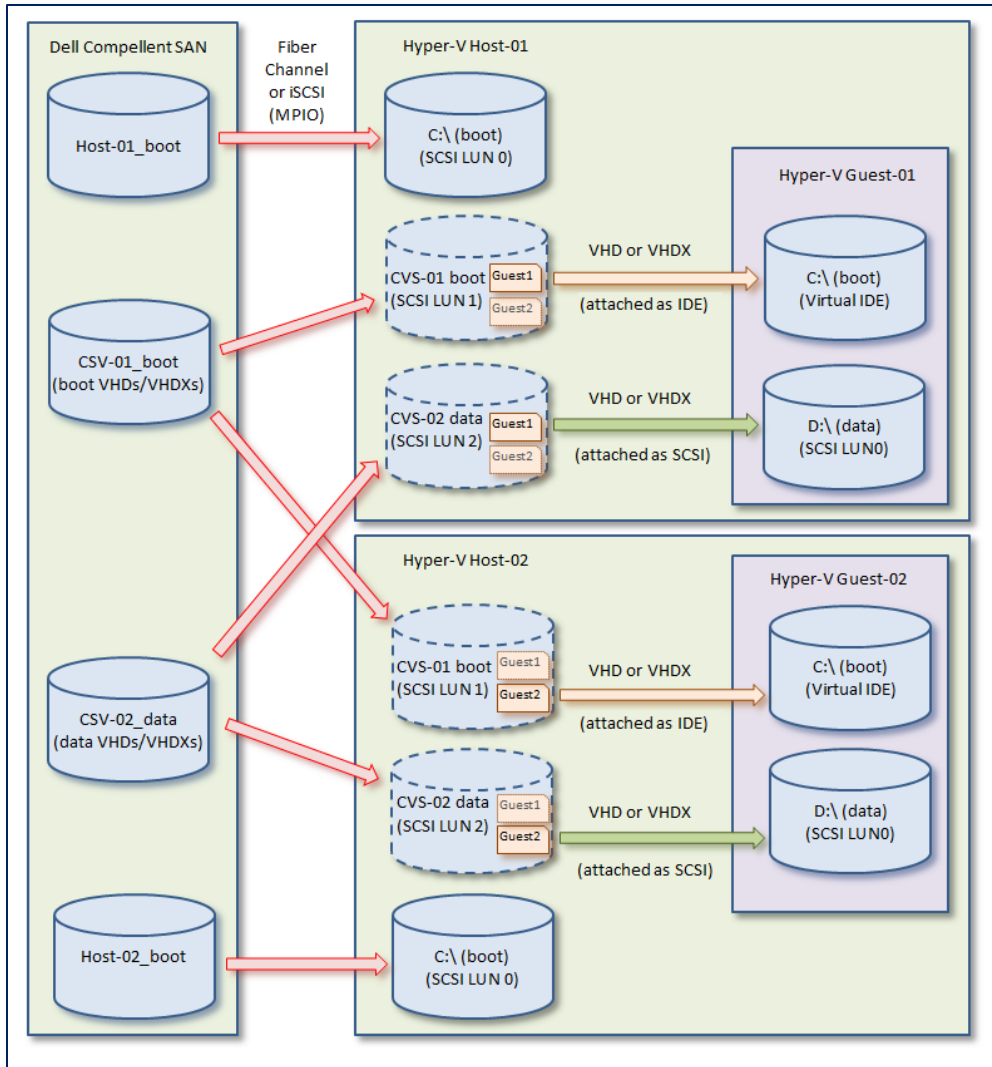


Figure 12. Hyper-V Guest VMs Configured to Use Cluster Shared Volumes (CSVs)

In the example shown in Figure 12, two CSVs (one for boot VHDs and another for data VHDs) are presented to two clustered Hyper-V nodes (Host-01 and Host-02). The two Hyper-V guests 01 and 02 each have a boot VHD or VHDX on CSV-01 and a data VHD or VHDX on CSV-02.

With this configuration, the boot VHDs/VHDXs on CSV-01 and the data VHDs/VHDXs on CSV-02 can be accessed concurrently by the two guest VMs even though the guest VMs are on two different nodes. The advantage of this configuration is either Guest-01 or Guest-02 can be live-migrated to the other node in the cluster without affecting the other guest VM, independently of which node has LUN ownership.

Note: For Hyper-V to function properly, the operating system of each node in the cluster must be configured to boot to the same drive letter. Normally, the boot drive will always be the C:\ drive on each server node. Additionally, The NTFS file system is required for all volumes enabled as Cluster Shared Volumes.

3.8.1 Advantages of Using Clustered Shared Volumes in a Failover Cluster

Using CSVs provides the following benefits in a failover cluster:

- Fewer LUNs are required for Hyper-V guests. Instead of having to manage one LUN per guest, many guests can share a single large LUN.
- Better SAN utilization.
- A guest VM can fail over to another node without requiring the other guests sharing the same CSV to also fail over.
- It is easier to manage the paths to VHD/VHDX files because they all reside under mount points under C:\ClusterStorage which are accessible by each cluster node.
- Cluster Validation for failover clusters runs quicker with fewer LUNs.
- There are no special hardware requirements (although CSVs must be formatted as NTFS volumes).
- Greater resiliency if there are SAN or network issues because the Cluster can redirect the CSV communication paths through parts of the SAN or network that may still be intact.

3.8.2 CSV Enhancements with Server 2012

- On Server 2008 R2, the Cluster Shared Volumes feature is only supported for use with Hyper-V. For Server 2012, CSV support has been extended beyond Hyper-V to include File Server, Storage Spaces and SMB 3.0.
- Support for ODX (offloaded data transfer).
- Better performance due to block-level (instead of file-level) redirection of IO.
- CSV's no longer have a dependency on Active Directory for authentication between cluster nodes.

4 Multipath I/O

The Windows Server operation system (2008 and newer) natively supports Multipath I/O (MPIO) by way of a Device Specific Module (DSM). MPIO is setup as a feature in the OS and must be enabled to work with the specific volumes that have multiple paths to the SAN. The process is very straightforward and simple to implement. The following procedures demonstrate how to install the feature, enable MPIO on Dell Compellent SAN volumes, and change the load balancing policy for specific volumes.

4.1 Install Windows Server Multipath I/O (MPIO)

The main purpose of Multipath I/O with Hyper-V (or any server for that matter) is to provide redundant paths for a server to access storage. With multipath, if one path goes down, another path is able to provide connectivity to prevent a service outage. MPIO also allows for load-balancing so that the I/O is spread across the available paths.

The process to install and configure Multipath MPIO is similar on a new install of either Window Server 2008 Hyper-V or Server 2012 Hyper-V. It is installed as an optional windows feature.

4.1.1 Install MPIO on a Windows 2008 or 2012 Server (with a GUI)

To set up MPIO on a Windows Server that has the GUI installed (not a core install), complete the following steps.

- 1) Enable a single path from Dell Compellent storage to an HBA on the server.
- 2) Stage the server OS using a single path to the storage.
- 3) Once the server OS has been installed, launch Server Manager.
- 4) Under Add Features, click the checkbox next to **Multipath I/O** as shown in Figure 13.

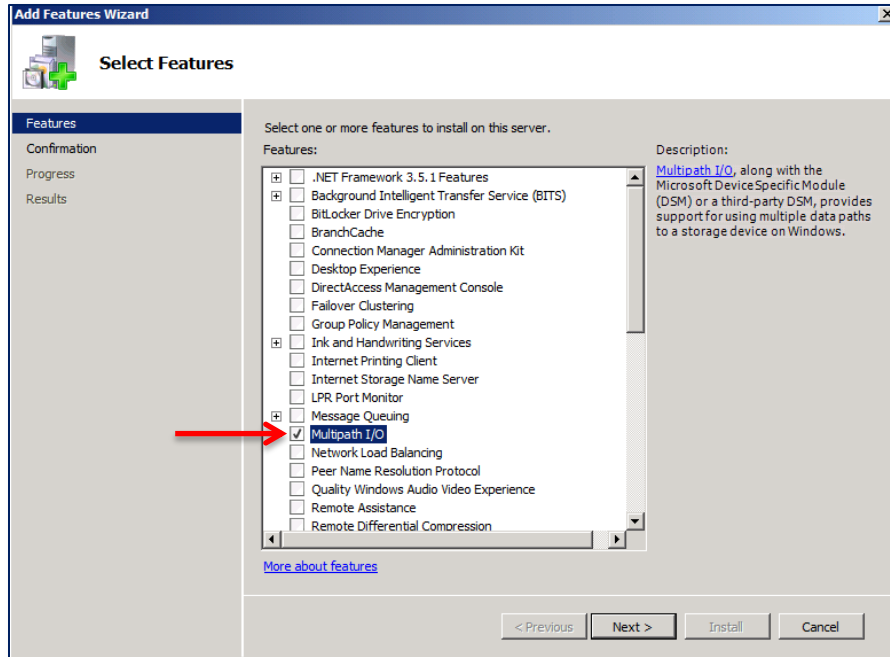


Figure 13. Install Multipath I/O (MPIO)

- 5) Allow the installation to complete.
- 6) Power down the server and using the BIOS settings for your HBA adapter, enable a second path to your Dell Compellent storage, and complete any required fabric or zoning changes.
- 7) Power up the server and log on as an administrator.

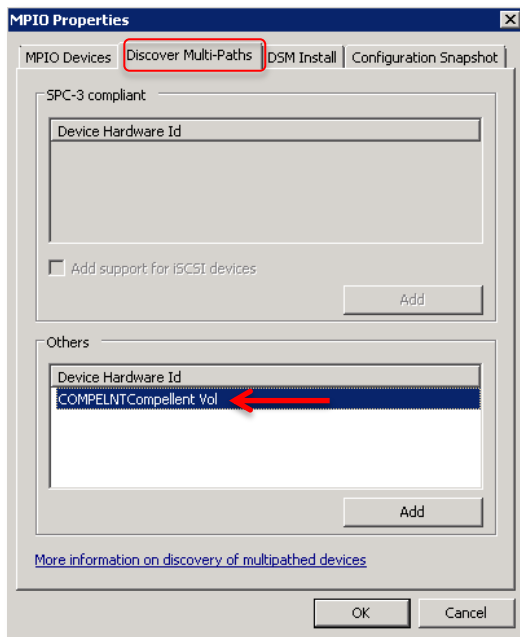


Figure 14. MPIO Devices – Discover Multi-Paths Tab

- 8) Under Administrative Tools, launch MPIO (or run **mpiocpl.exe** from a command

prompt).

- 9) Click on the **Discover Multi-Paths** Tab.
- 10) Click on the **COMPELNTCompellent Vol** line in the **Others** box, then click on the **Add** button.

Note: If **COMPELNTCompellent Vol** is not listed under the **Discover Multi-Paths** tab, make sure that at least two paths are enabled for the server to access the SAN storage. If the server can only see one path, **COMPELNTCompellent Vol** will not be listed.

- 11) After clicking on the **Add** button, the server will prompt for a reboot. Reboot the server.
- 12) Once the server has rebooted, under Administrative Tools, launch MPIO (or run **mpiocpl.exe** from a command prompt).

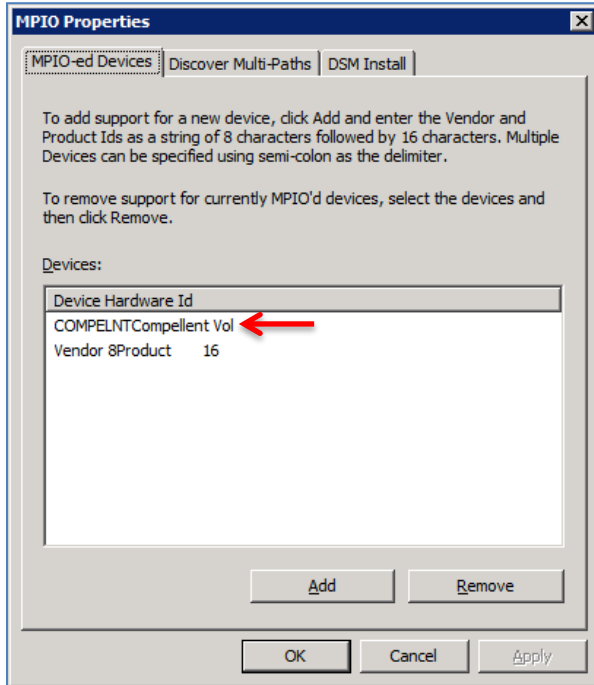


Figure 15. MPIO Properties Screen Displays MPIO Devices

- 13) Verify that **COMPELNTCompellent Vol** is listed under the **MPIO** tab as shown in Figure 15.

4.1.2 Install MPIO on a Server 2008 Core Installation

To install MPIO for a Server 2008 core installation, run this command from a command prompt:

```
ocsetup MultipathIo /norestart
```

Next, run:


```
mpclaim -r -i -a ""
```

Usage: mpclaim <reboot_option> <install_switch> <device_switch> <device_hwid(s)>

- <reboot_option> - whether or not to auto reboot.
 - -r automatically reboot without prompting.
 - -n suppress reboot request.
- <install_switch> - whether to add or remove MPIO support.
 - -i install MPIO optional component and add multipath support for device.
 - -u remove multipath support for device and uninstall MPIO.
- <device_switch> - whether to apply above options to all devices or passed-in devices
 - -d following parameters indicate hardware ids of devices
 - -a work on all applicable devices
 - -c work on only all SPC3-compliant devices (meaningful only in the context of install. If used with '-u', it is treated as '-a')
- <"device_hwid"> - hardware IDs of devices to be MPIO-ed, as strings of vendor8product16, delimited by a space (Note: Use empty string with '-a' option)

The command syntax as shown above will install the Microsoft MPIO optional component on server 2008 core installations and then set the Microsoft MPIO DSM (MSDSM) to claim all unclaimed Microsoft MPIO devices in the system.

MPIOCPL.exe is included in Server Core and can be launched from a command line to bring up the GUI to configure MPIO devices:

```
c:\>MPIOCPL.EXE
```

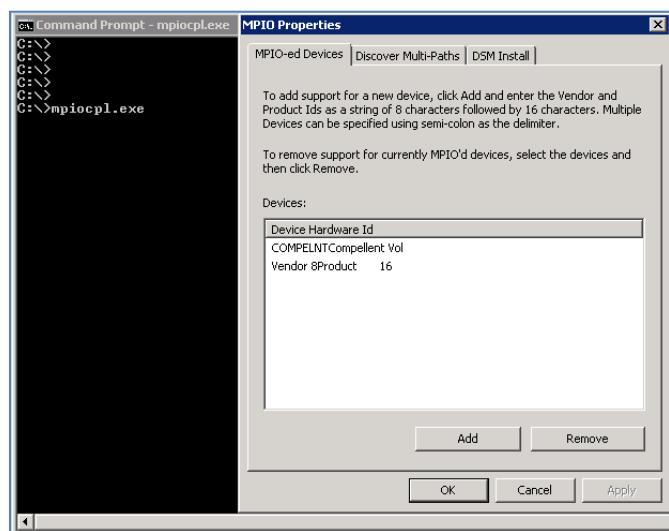


Figure 16. MPIOCPL Configuration Screen

4.1.3 Install MPIO on a 2012 Server Core Installation

To install the MPIO feature on a Windows 2102 Server with a core OS install (no GUI) complete the following steps:

- 1) Enable a single path from Dell Compellent storage to an HBA on the server.
- 2) Stage the server OS using a single path to the storage.
- 3) Using a command window or PowerShell, configure the server's name, domain, IP, DNS, firewall, and other basic settings as desired.
- 4) Because a Windows 2012 Server with a core OS install does not have a GUI, it is not possible to run Server Manager from the server's console. However, it is very easy to do so remotely from another Windows 2012 Server that has a GUI.
- 5) Launch Server Manager from another Windows 2012 Server that has a GUI.

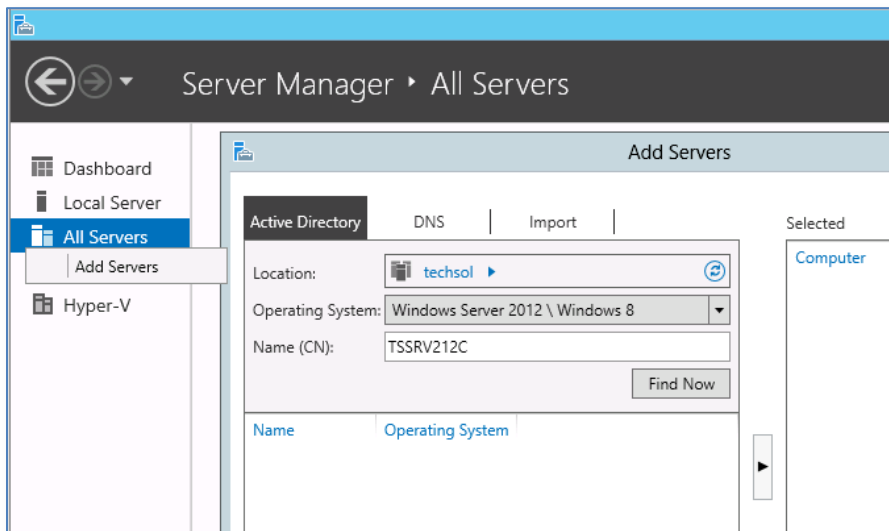


Figure 17. Add a Remote Server to Server Manager

- 6) Right click on **All Servers** and select **Add Servers** to launch the Add Servers dialog window. Complete the steps to find and select the desired Windows 2012 core server for remote management.

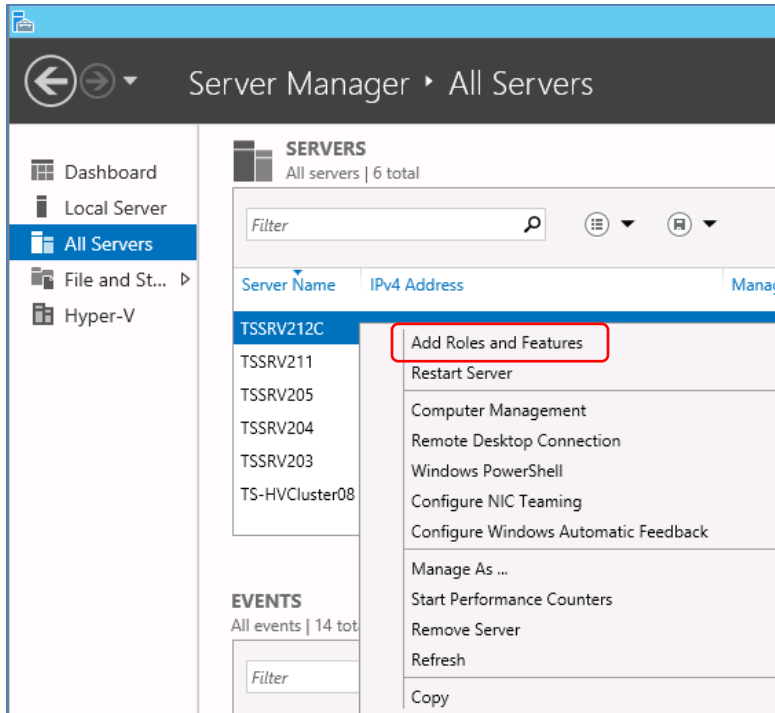


Figure 18. Centrally Manage Roles and Features Using Server Manager

7) Once the core server has been added to Server Manager, right click on it and select **Add Roles and Features** as shown in Figure 18 to launch the wizard.

8) Under Add Features, click the checkbox next to **Multipath I/O** as shown in Figure 19.

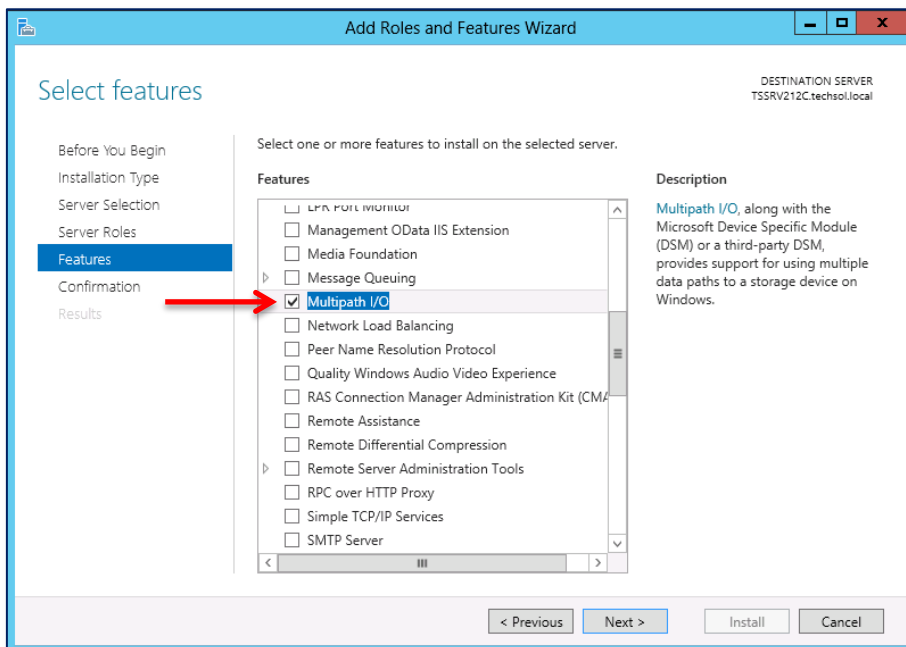


Figure 19. Install Multipath I/O (MPIO) (Windows 2012 with Core OS Install)

- 9) Allow the installation to complete.
- 10) Power down the server and using the BIOS settings for your HBA adapter, enable a second path to your Dell Compellent storage, and complete any required fabric or zoning changes.

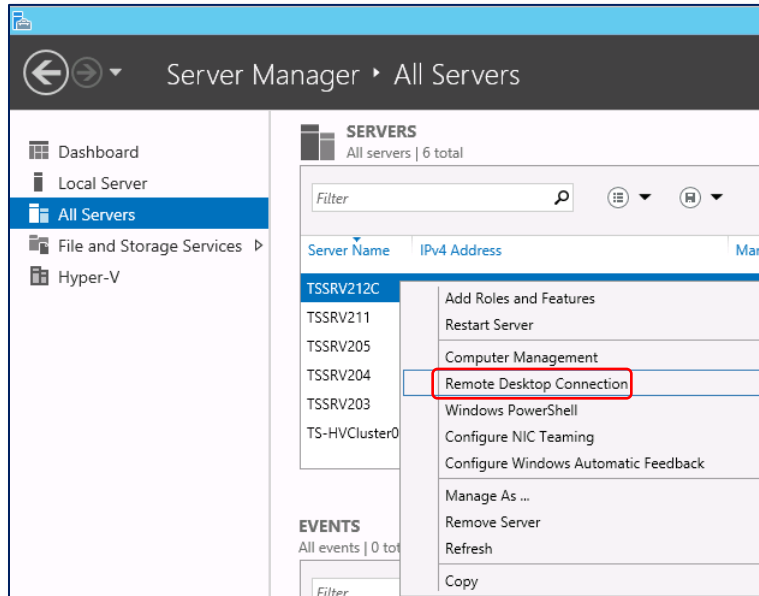


Figure 20. Remote Desktop Connection to a Core Server from Server Manager

- 11) Power up the core server and use Remote Desktop Connection from Server Manager to access the core server.

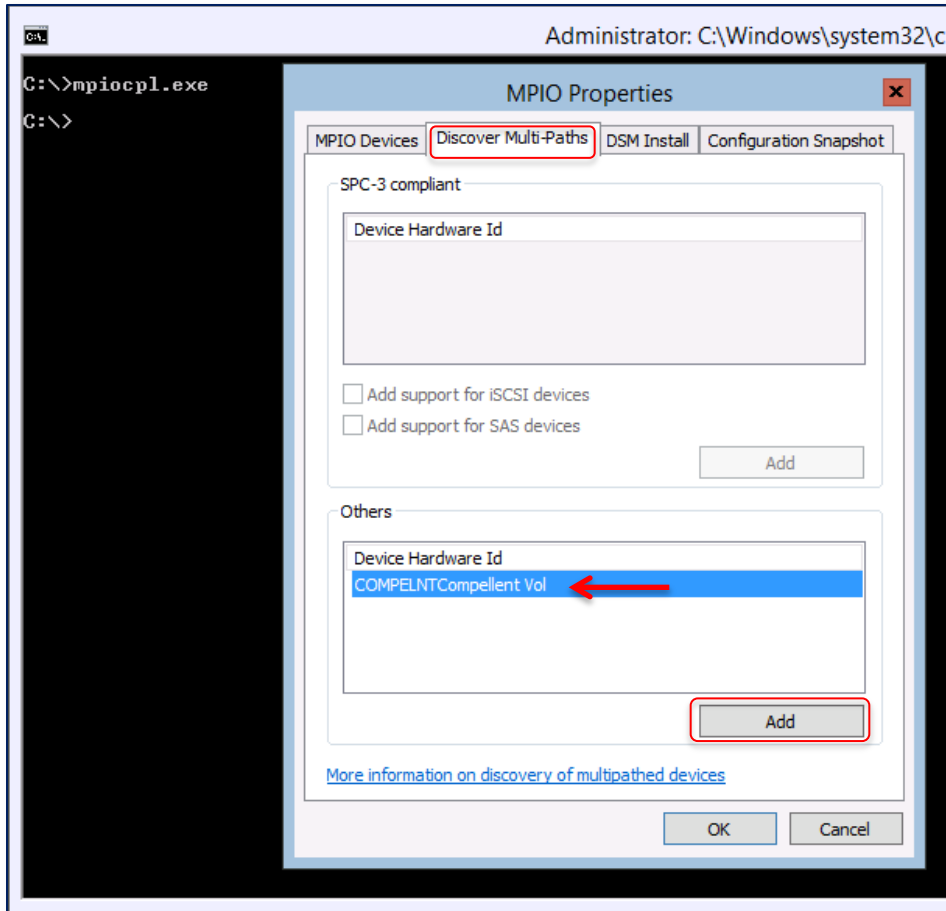


Figure 21. MPIO Devices – Discover Multi-Paths Tab

- 12) At the command prompt on the core server type **mpioctrl.exe** to launch the MPIO Properties window.
- 13) Click on the **Discover Multi-Paths** Tab.
- 14) Click on the **COMPELNTCompellent Vol** line in the **Others** box, then click on the **Add** button.

Note: If **COMPELNTCompellent Vol** is not listed under the **Discovered Multi-Paths** tab, make sure that at least two paths are enabled for the 2012 core server to access the SAN storage. If the 2012 core server can only see one path, **COMPELNTCompellent Vol** will not be listed.

- 15) After clicking on the **Add** button, the server will prompt for a reboot. Reboot the server.
- 16) Once the server has rebooted, use Remote Desktop from Server Manager to connect to the 2012 core server.
- 17) At the command prompt, run **mpioctrl.exe** to launch the MPIO Properties window.

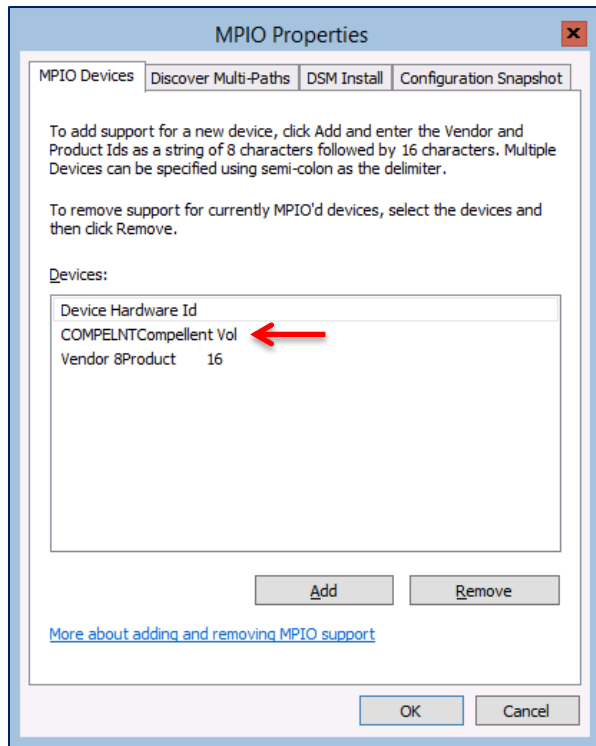


Figure 22. MPIO Properties Screen Displays MPIO Devices

18) Verify that **COMPELNTCompellent Vol** is now listed under the **MPIO Devices** tab as shown in Figure 22. The default load-balancing policy is set to round robin.

4.2 Configure MPIO Load Balancing Policy

By default the MPIO load balancing policy for Windows 2008 R2 and Server 2012 is set to **Round Robin**. For Windows 2008 (prior to R2) the default is set to **Fail Over Only**. The load balancing policy is set for (and can be changed for) each individual volume connected to the Dell Compellent SAN (for which multiple paths exist).

Note: For Windows Server 2008 (prior to R2) the Microsoft MPIO load balancing policy can only be configured in Windows 2008 full installations. Windows 2008 Server Core installations (prior to R2) cannot be changed from the default MPIO policy of **Fail Over Only**.

The load balancing policy can be changed per volume to one of the following five options:

- Fail Over Only: employs one active/optimized path and all other paths are active/unoptimized, standby, or unavailable. The non-active/unoptimized paths will be tried round-robin upon failure of the active/optimized path until an available path is found.
- Round Robin: attempts to evenly distribute incoming requests to all processing paths.
- Round Robin with subset: This load balancing policy is not supported by Dell Compellent.

- Least Queue Depth: compensates for uneven loads by distributing I/O proportionately across paths.
- Weighted Paths: allows the administrator to specify the relative processing load of each path. A higher number indicates a lower priority path.

4.2.1 Change the Load Balancing Policy on a Volume

This must be completed for each volume that is added to the system if a different policy than the default is desired.

- 1) Access the Computer Management console.
 - a. For Server 2008, click on **Start→Administrative Tools→Computer Management**.
 - b. For Server 2012, launch Server Manager from the task bar, and then click on **Tools→Computer Manager**.

Note: For Core installs of Server 2008 R2 and Server 2012, please see Section 4.2.2 to configure load balancing via the command line or PowerShell.

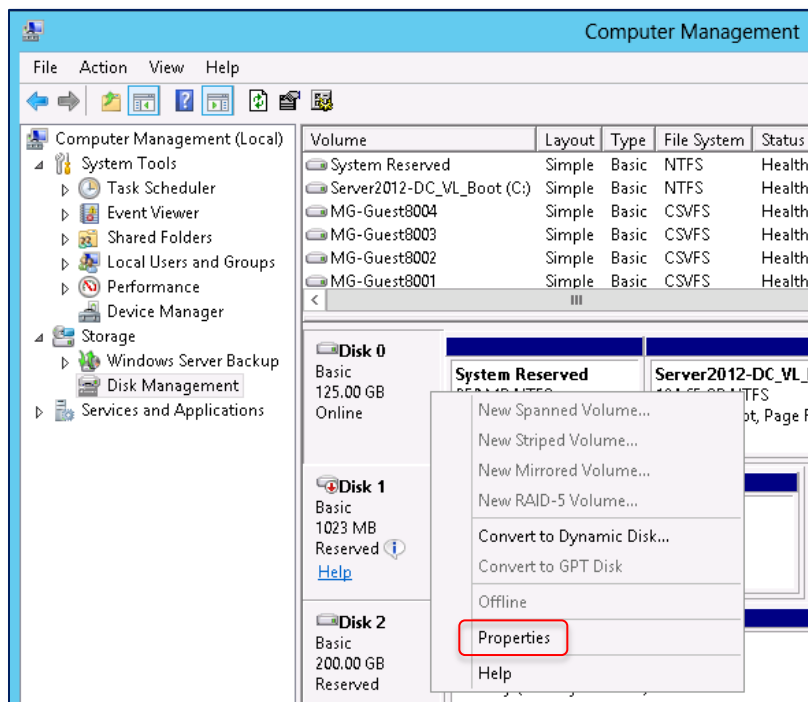


Figure 23. Disk Management Disk Properties

- 2) Navigate to the desired volume under **Storage→Disk Management**.
- 3) Right click on the desired volume and choose **Properties**.




Figure 24. Select MPIO Options for a Volume

- 4) On the **Properties** screen, select the **MPIO** tab.
- 5) From the drop down list, select the desired MPIO policy for the environment being configured, and then click on OK.
- 6) Repeat steps 3 – 5 for each additional volume that is configured to use MPIO.
- 7) As a time saver, if a large number of volumes on a server need to have their load balance policy changed to the same settings, or if a different default policy for new volumes is desired, use **MPClaim** from the command prompt (see Section 4.2.2).

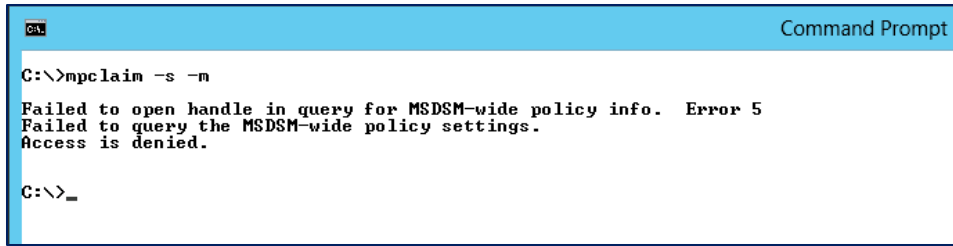
Note: If a system-wide policy has not been set by using MPClaim, new volumes added to a server will automatically use the Windows OS default MPIO setting (e.g. Round Robin for Server 2008 R2 and Server 2012). To change the default policy to something else, use MPClaim (see Section 4.2.2) to change the Windows OS default. Any new volumes added to the server will then use the specified policy.

4.2.2 Change the Load Balancing Policy on Volumes by Using MPClaim

MPClaim.exe can be used from a command prompt to change the load balancing policy for:

- Windows 2008 Full (but not Core) installations
- Windows 2008 R2 Full and Core installations
- Server 2012 Full and Core Installations.

Note: While MPClaim is supported with Server 2012, Microsoft recommends using the PowerShell Cmdlets included with Windows 2012 to change load balancing policies.

A screenshot of a Windows Command Prompt window. The title bar reads "Command Prompt". The command prompt shows the command `C:\>mpclaim -s -n` followed by an error message: `Failed to open handle in query for MSDSM-wide policy info. Error 5`, `Failed to query the MSDSM-wide policy settings.`, and `Access is denied.` The prompt then shows `C:\>_` indicating the command has finished.

```
C:\>mpclaim -s -n
Failed to open handle in query for MSDSM-wide policy info. Error 5
Failed to query the MSDSM-wide policy settings.
Access is denied.
C:\>_
```

Figure 25. Run Command Prompt with Administrator Rights to Avoid Access Denied Errors

- 1) Open a command prompt window with full administrator rights. If the error displayed in Figure 25 is experienced, close the command prompt window, and re-launch it by right-clicking on the command prompt icon and choosing **Run as Administrator**.
- 2) To show the current target-wide load balance policy, type the following:

```
mpclaim -s -t
```

- 3) To show the current default MSDSM-wide load balance policy, type:

```
mpclaim -s -m
```

- 4) As an example, to set the default MSDSM-wide load balance policy of all LUNs controlled by the Microsoft DSM to **Least Queue Depth** type:

```
mpclaim -l -m 4
```

- 5) For more information on **MPClaim.exe** options and switches, type **mpclaim.exe /?** from the command prompt.

4.2.3 Configuring Server 2012 MPIO Settings with PowerShell

New to Server 2012 is the ability to enable or disable MPIO, or change MPIO settings by using the PowerShell Cmdlets included with the Server 2012 OS.

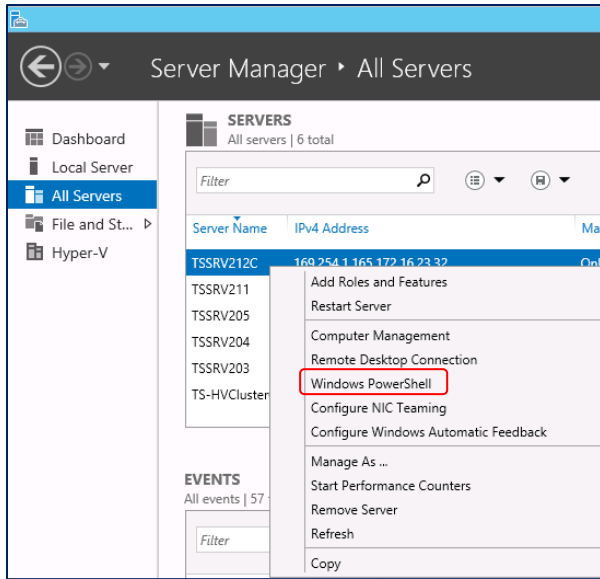


Figure 26. Use Server Manager to Launch PowerShell for a Remote 2012 Server Core

- 1) Launch a PowerShell command window with elevated (administrator) privileges.
 - If managing a 2012 Server remotely using Server Manager, right click on the server to be managed and select Windows PowerShell as shown in Figure 26.
 - Remote PowerShell access to a server works for either Full or Core installs of Server 2012.
- 2) To verify that the MPIO feature is installed, type:

```
Get-WindowsOptionalFeature -Online -FeatureName MultipathIO
```

- 3) To enable MPIO, type:

```
Enable-WindowsOptionalFeature -Online -FeatureName MultipathIO
```

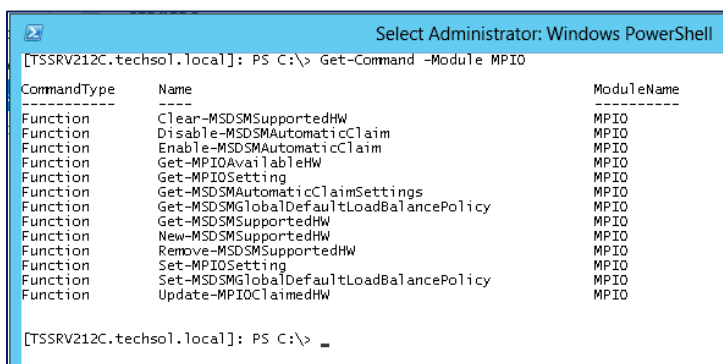


Figure 27. Server 2012 MPIO PowerShell Commands

- 4) To list the commands available for the MPIO module, type:

```
Get-Command -Module MPIO
```

For more information on using PowerShell to set MPIO setting with Server 2012, please refer to the following documentation (also listed in the References section):

Microsoft Multipath I/O (MPIO) Users Guide for Windows Server 2012

<http://www.microsoft.com/en-us/download/details.aspx?id=30450>

4.3 iSCSI Software Initiator MPIO Configuration

4.3.1 iSCSI Overview

Server 2008 R2 and 2012 Hyper-V hosts and guest VMs support MPIO with the iSCSI software initiator. iSCSI initiator software is provided natively by the Windows Server OS.

The iSCSI “quick connect” feature works well for single iSCSI path connectivity. Configuring iSCSI to use MPIO requires a few more steps, but is still easy to configure.

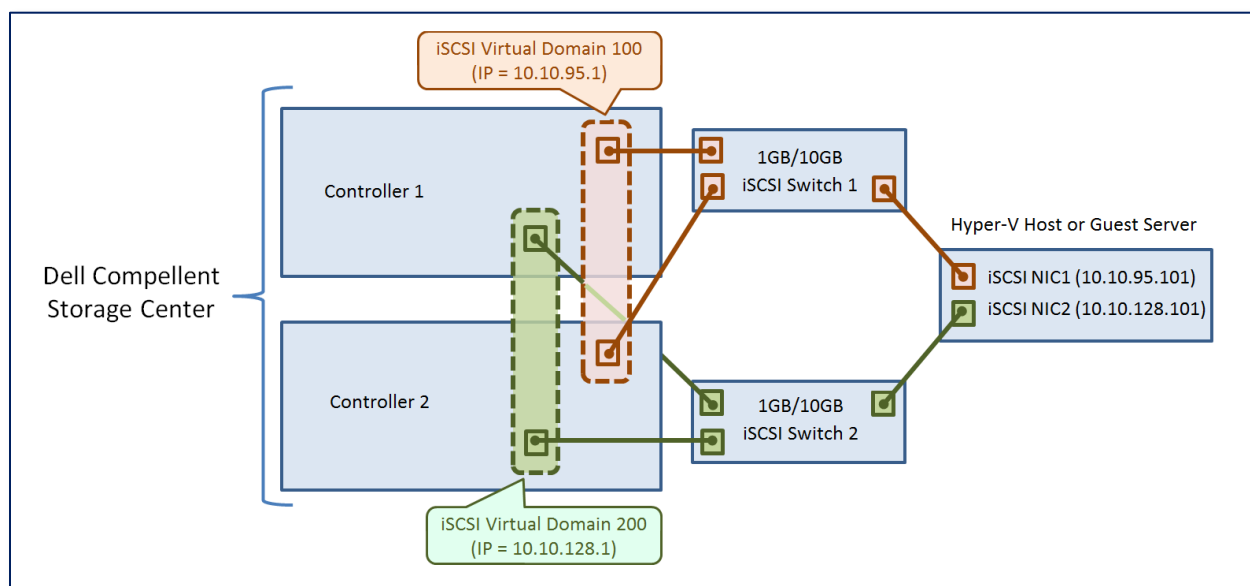


Figure 28. iSCSI Virtual Ports and Domains

Figure 28 represents a dual-controller Dell Compellent Storage Center that is configured with virtual front-end ports with 2 fault domains. Two physical iSCSI ports (one from each controller) are grouped logically as a virtual domain that is assigned a virtual iSCSI IP address. Each virtual domain’s physical ports are connected to two separate iSCSI switches to ensure full path redundancy to the dual iSCSI NICs on the host and guest VMs.

4.3.2 Configure a Server 2008 R2/2012 Hyper-V Host for iSCSI MPIO

Environment overview (refer to Figure 28):

- Windows Server 2008 R2 Hyper-V or Server 2012 Hyper-V server host.
- Two dedicated NIC ports (one for each fault domain) on the host server for iSCSI

communications.

- A Dell Compellent Storage Center with dual controllers configured with virtual front-end iSCSI ports and two virtual fault domains.
- The MPIO feature is installed on the server host.

- 1) Configure two physical NIC ports on your Hyper-V host to use iSCSI.
 - a. Configure one NIC (along with any necessary cabling and switching configurations) to access one virtual iSCSI domain, and configure the 2nd NIC to access the other virtual iSCSI domain.
 - b. In this example as show in Figure 28, the host server's NICs are assigned with IPs of 10.10.95.101 (for fault domain 100), and 10.10.128.101 (for fault domain 200).
 - c. Once configured, the Hyper-V host should be able to ping both of the virtual iSCSI IP addresses associated with the two fault domains on the Dell Compellent Storage center, in this example, 10.10.95.1 (fault domain 100) and 10.10.128.1 (fault domain 200).
- 2) Logon to the Windows 2008 R2 or 2012 Hyper-V host server and launch the iSCSI Initiator software.

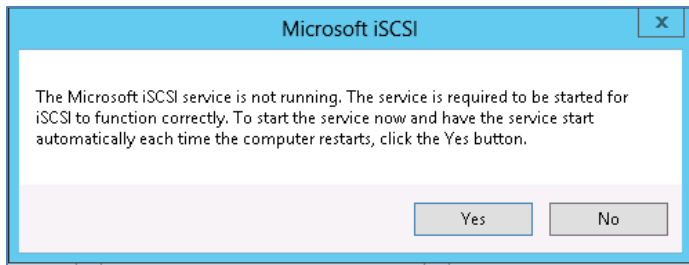


Figure 29. Start the Microsoft iSCSI Service Prompt

- a. For Server 2008, go to **Start→Administrative Tools→iSCSI Initiator**.
- b. For Server 2012, start Server Manager, and from the Dashboard, click on **Tools→iSCSI Initiator** from the drop-down list.

Note: If receiving a prompt to start the iSCSI service (for either Server 2008 R2 or Server 2012), click on Yes.

- 3) Select the **Discovery** tab then click the **Discovery Portal** button.

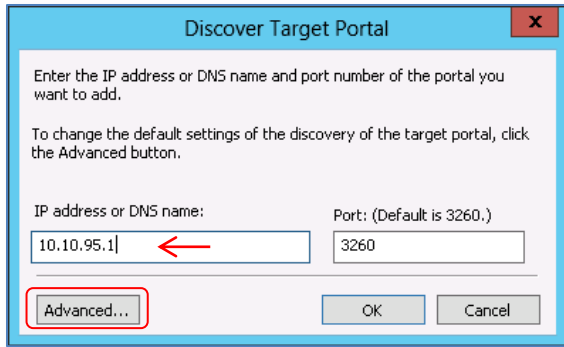


Figure 30. Enter iSCSI Target IP

- 4) Enter the IP address of the first iSCSI target IP, in this example, the virtual iSCSI IP addressed associated with the first virtual fault domain on the Dell Compellent Storage Center (see Figure 28), and then click on the **Advanced** Button.

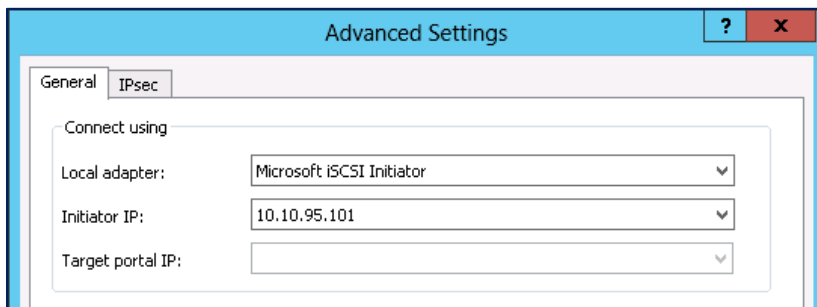


Figure 31. Advanced Settings for iSCSI Initiator

- 5) From the **Local adapter** drop-down list, select **Microsoft iSCSI Initiator**.
- 6) From the **Initiator IP** drop-down list, select the local IP address of the Hyper-V host server's NIC that is to be associated with the first fault domain (fault domain 100) which in this example is 10.10.95.101.

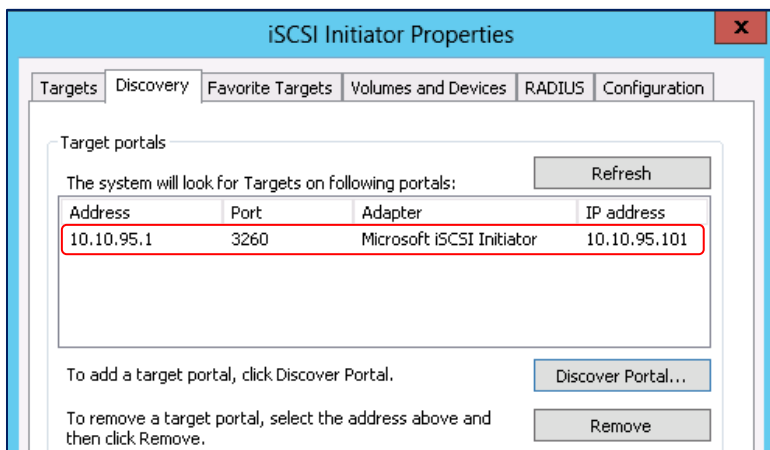


Figure 32. Target Portal for First Fault Domain is Now Listed

- 7) Click **OK**, and then **OK** again to return to the iSCSI Initiator properties window.
- 8) Verify that the target IP address and adapter IP address now show in the Target Portals window as shown in Figure 32.
- 9) Repeat steps 3 – 8 above to add the 2nd target IP for the 2nd virtual fault domain and the Hyper-V host's 2nd iSCSI NIC (in this example, 10.10.128.1 and 10.10.128.101)

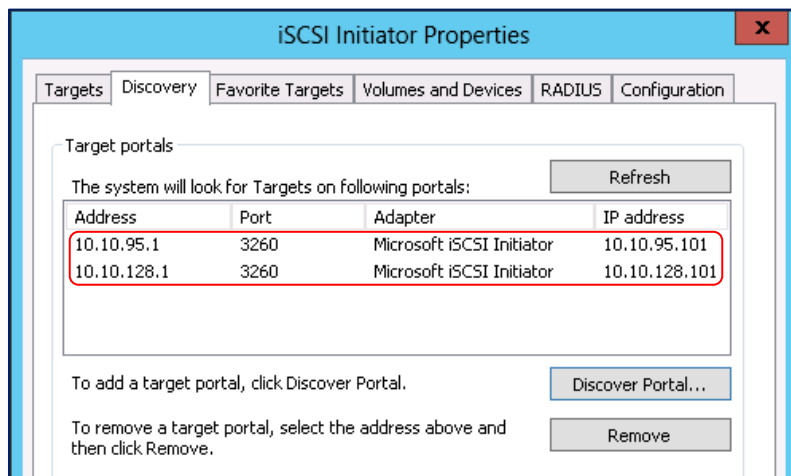


Figure 33. Both Sets of Target and Initiator IPs Are Now Listed

- 10) When completed, both pairs of iSCSI initiators and targets should be listed as shown in Figure 33.

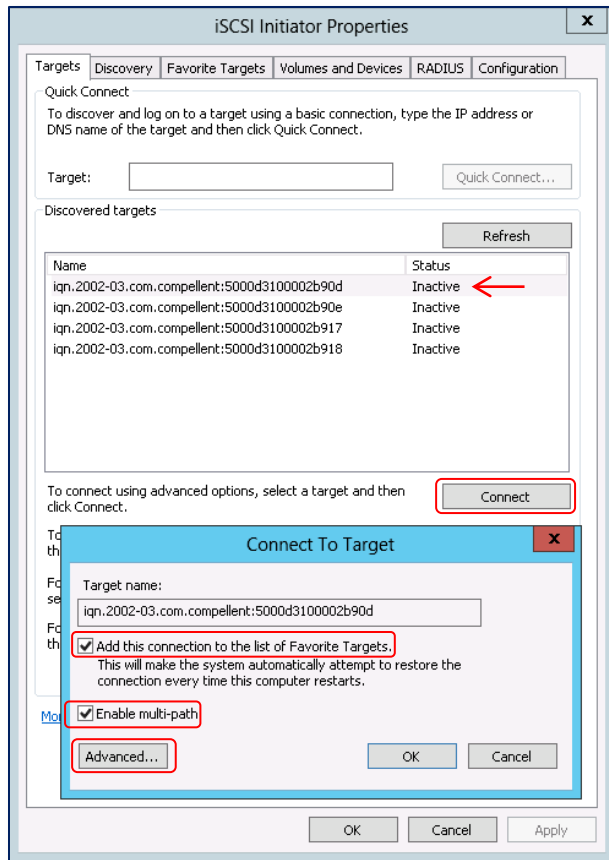


Figure 34. Activate Targets and Enable Multipath

- 11) Select the **Targets** Tab. This should be populated with the discovered iSCSI target ports on the Dell Compellent Storage Center.
- 12) Highlight the first target by clicking on it, and then click on the **Connect** button.
- 13) On the **Connect To Target** screen, verify that both **Add this connection to the list of Favorite Targets** and **Enable multi-path** are checked.
- 14) Click on the **Advanced** button.

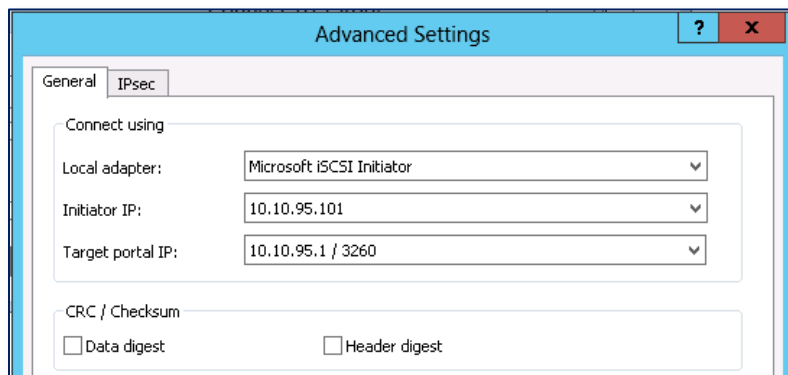


Figure 35. Advanced Settings for iSCSI Target Connections

- 15) Set **Local adapter** to **Microsoft iSCSI Initiator**.

- 16) Select the **Target portal IP** from the dropdown list.
- 17) Select the **Initiator IP** from the dropdown list.
- 18) Then click **OK**, and then **OK** again to return to the iSCSI Initiator properties window.
- 19) Repeat steps 12 - 18 above for each additional target listed.

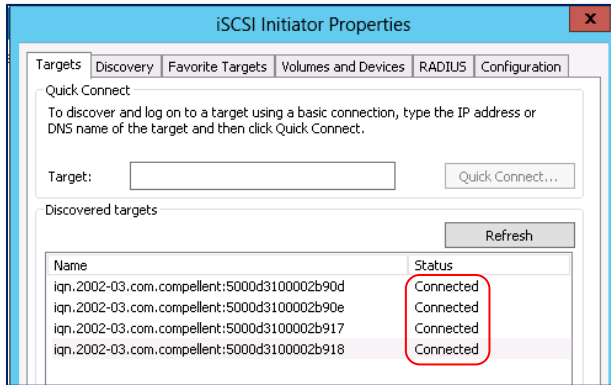


Figure 36. Targets Now Show as Connected

- 20) When finished, all the targets should show as connected as shown in Figure 36.
- 21) Click on **OK** to exit the iSCSI Initiator Properties window.

4.3.3 Configure the Dell Compellent Storage Center

- 1) Log in to the Dell Compellent Storage Center Manager GUI.
- 2) Expand the **Servers** tree and select (or create) the desired server folder for the server object.
- 3) Right click on the desired server folder and select **Create Server**.

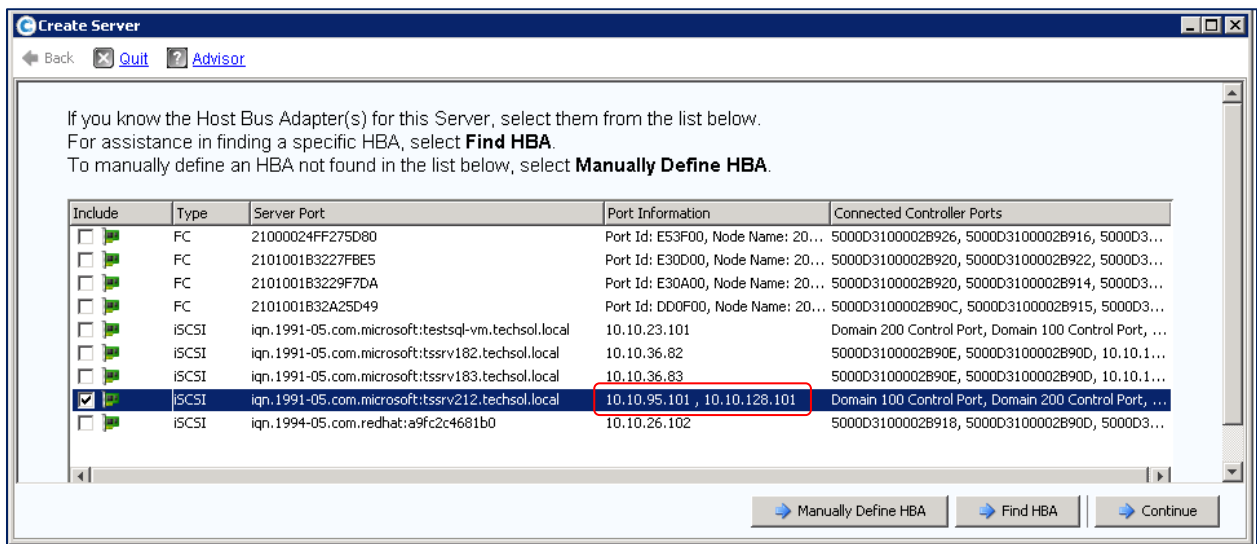


Figure 37. Select iSCSI HBA

- 4) Select the iSCSI server port associated with the iSCSI IPs assigned to the server, in this example, 10.10.95.101 and 10.10.128.101 as shown in Figure 37, then click on **Continue**.

- 5) On the next screen, verify the server folder for the server object, provide a name for the server object, indicate the operating system (with MPIO), and then click on **Continue** and **Create Now**.



Figure 38. Create Volume for New Server

- 6) To create a volume now and add assign it to the server, click on **Create Volume**. To create a volume later, click on **Close**.

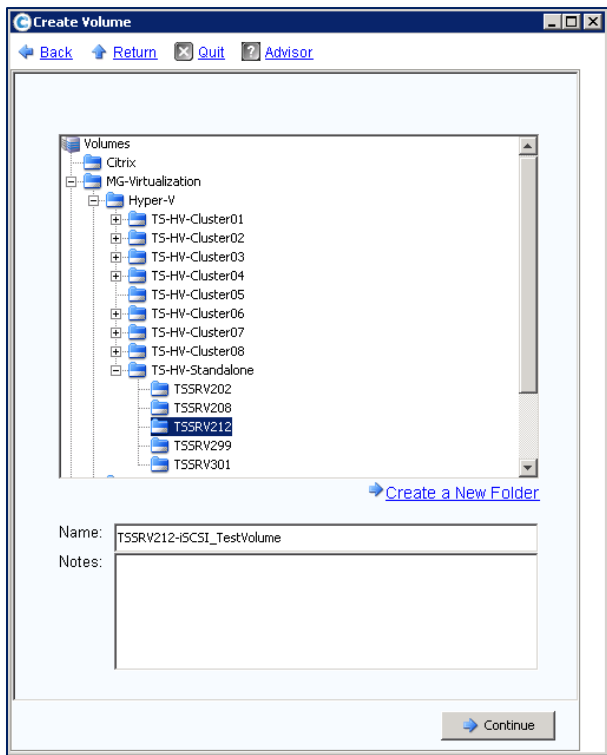


Figure 39. Create Volume for New Server

- 7) After clicking on **Create Volume**, specify the volume size, volume name, volume

- folder location, click on **Continue**, then on **Create Now**.
- 8) On the next screen, select **Map this Volume to this Server**.
 - 9) Click on the **Advanced** button to review any advanced settings, then click on **Create Now**, and then on **Close**.

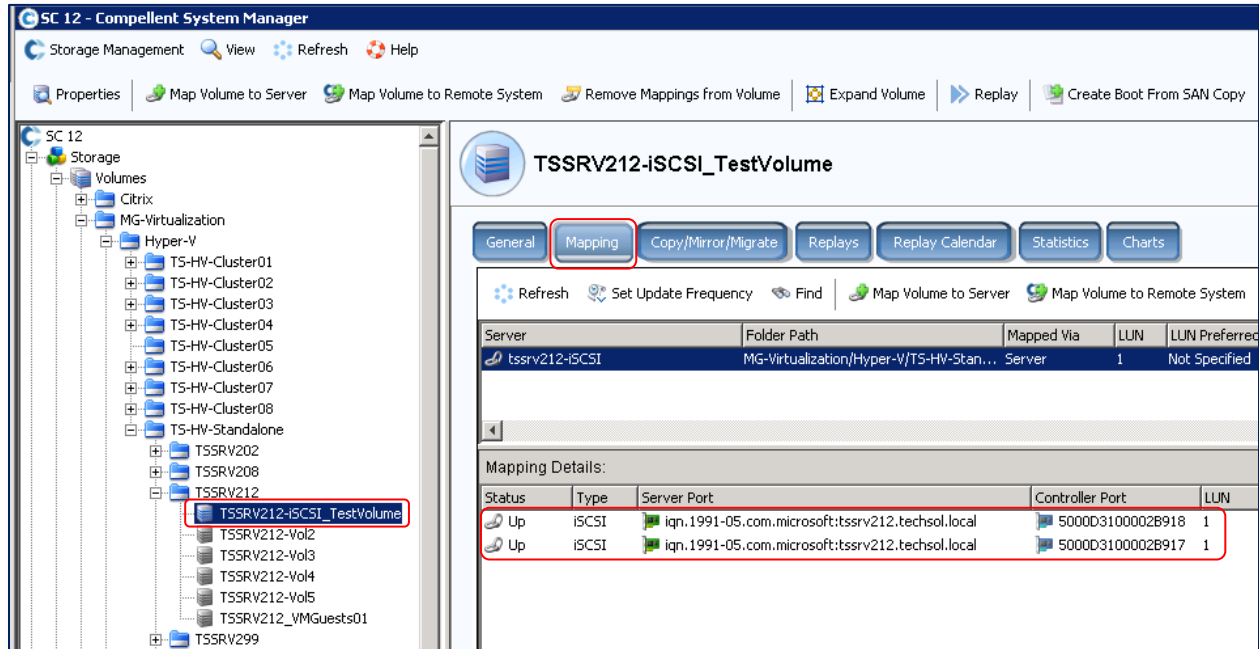


Figure 40. View the iSCSI Mappings for the New Volume

- 10) Highlight the new volume, and click on the **Mappings** tab. Note that two iSCSI ports are listed.

4.3.4 Configure MPIO for the iSCSI Volume on the Hyper-V Server

- 1) Logon to the Hyper-V host server and access Disk Manager
 - a. For Server 2008, click on **Start→Administrative Tools→Server Manager**, and then expand **Storage** and click on **Disk Management**.
 - b. For Server 2012, start **Server Manager** from the task bar and select **Computer Manager** from the **Tools** drop-down list, then expand **Storage** and click on **Disk Management**.
- 2) From the **Action** drop down menu, click on **Rescan Disks**.

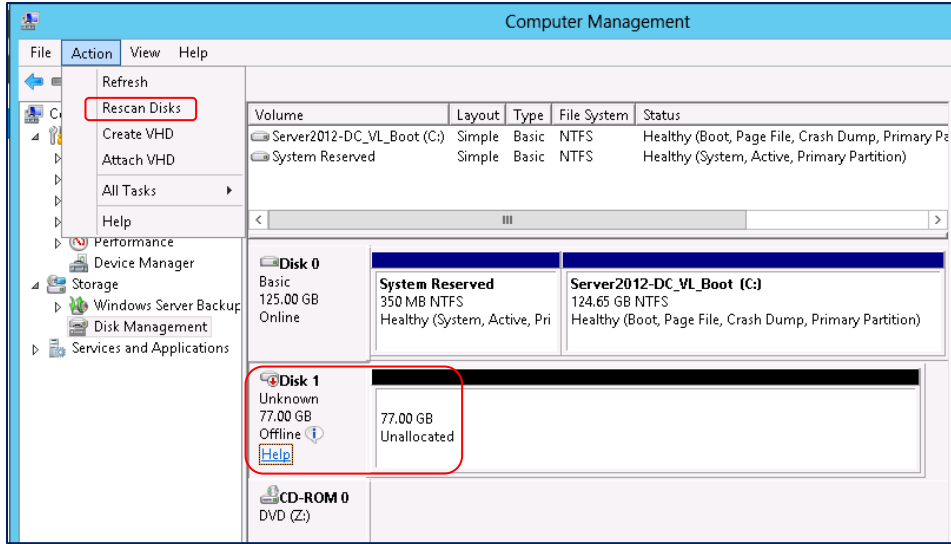


Figure 41. User Disk Management to Rescan Disks

- 3) Verify that the new volume shows up in the list, and then minimize Server Manager or Computer Management.
- 4) Launch the iSCSI Initiator software.
 - a. For Server 2008, go to **Start**→**Administrative Tools**→**iSCSI Initiator**.
 - b. For Server 2012, start Server Manager, and from the Dashboard, click on **Tools**→**iSCSI Initiator** from the drop-down list.

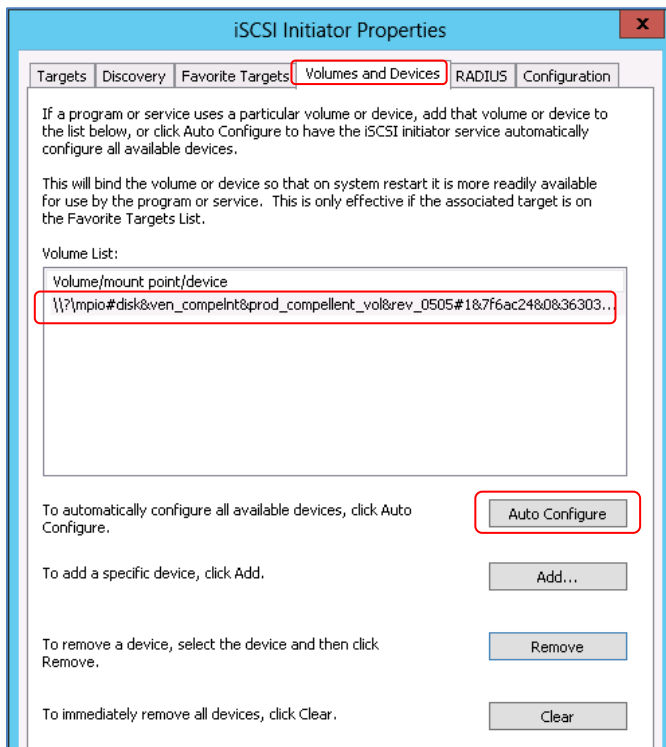


Figure 42. iSCSI Volumes and Devices

- 5) Click on the **Volumes and Devices** tab.

- 6) If the volume is not already listed, select the **Auto Configure** button.
- 7) Once the volume is listed, then close out of the iSCSI Initiator window by clicking on **OK**.
- 8) Under **Administrative Tools**, launch MPIO (or run **mpiocpl.exe** from a command prompt).

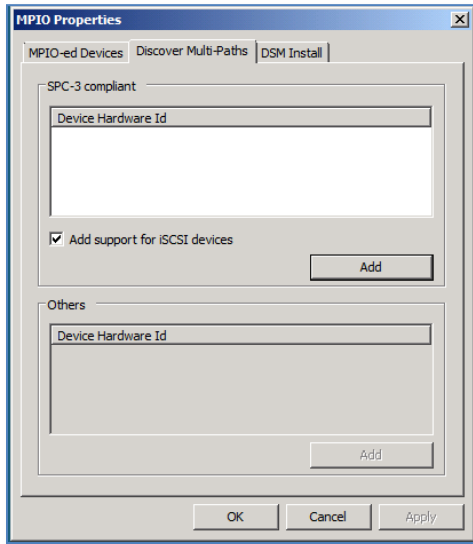


Figure 43. Discover Multi-Paths

- 9) Select the **Discover Multi-Paths** tab as shown in Figure 43.
- 10) Check **Add support for iSCSI devices** then click the **Add** button.

Note: If the **COMPELNTCompellent Vol MSDSM** is already installed (listed under the MPIO Devices Tab), then the checkbox to **Add support for iSCSI devices** will be grayed out. This is by design. If the MSDSM is already installed, no further action is necessary – exit out of MPIO Properties, and use Disk Manager to set the MPIO policy for the new volume.

- 11) After checking the **Add Support for iSCSI Devices**, reboot the server when prompted.

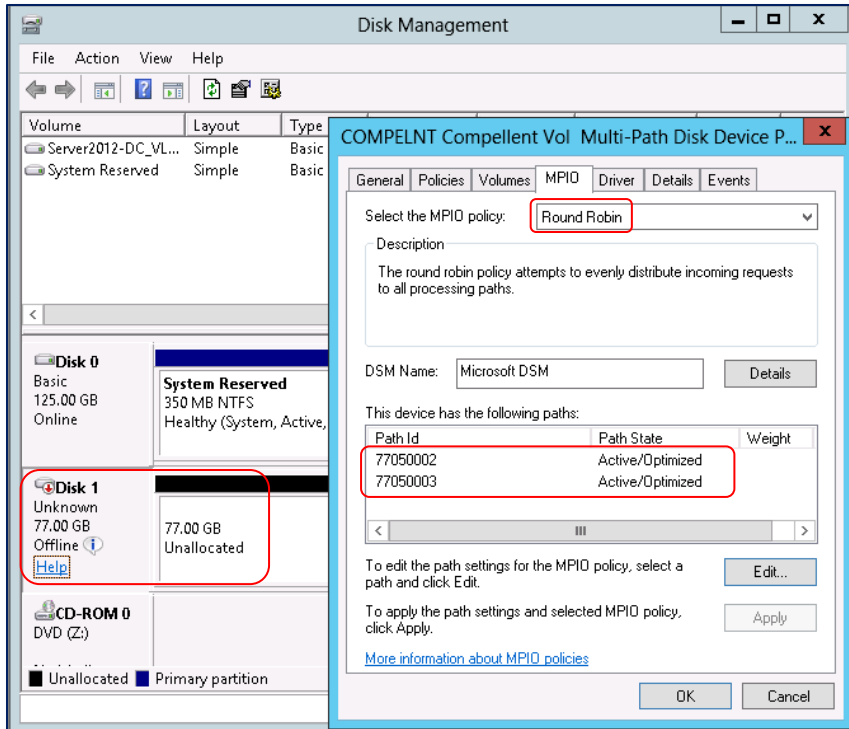


Figure 44. Configure MPIO Settings for the iSCSI Volume

- 12) Once the server has rebooted, MPIO now should be enabled on the volume. Set the MPIO load balance policy on the volume as desired (use MPClaim.exe from a command prompt if desired).

5 Data Instant Replay

The Dell Compellent Storage Center allows for the creation of Data Instant Replays (snapshots) that enable administrators to do the following:

- Recover servers to crash-consistent states including Hyper-V host servers and guest VMs.
- Quickly spin up lab or testing environments from Replays of production servers or data.
- Create gold images and use shared View Volumes to save disk space when adding additional Hyper-V servers and guests VMs.

Dell Compellent Replays can be taken of disks with VHDs/VHDXs, pass-through disks, and directly-attached iSCSI or virtual fiber channel volumes mapped to host servers or guest VMs.

Dell Compellent Storage Center Replays are “crash-consistent” snapshots. When using a crash consistent Replay to recover a server, it is similar to having the server recover from a power outage from the point in time the Replay was taken. In most cases, servers can be recovered from crash consistent Replays without any issues.

If additional server protection is desired by capturing application-consistent VSS-integrated Replays of Hyper-V guest VMs, please refer to the Dell Compellent Replay Manager 6 Users Guide. Dell Compellent Replay Manager 6 is able to leverage Microsoft VSS to take application-consistent (IO is paused) Replays of Hyper-V guests, Exchange servers, and SQL servers.

The following section provides examples for how to configure and use Data Instant Replay within a Hyper-V environment.

5.1 Using Replays to Recover Guest VMs

Hyper-V guest VMs can be recovered by using Dell Compellent Replays. The easiest and quickest method of recovery is when there is a one-to-one configuration (a single Hyper-V guest VM per LUN). Simply create a View Volume from a Replay of the LUN and present the View Volume to the Hyper-V host as a new LUN, and the guest VM can be recovered to that point in time.

Individual Hyper-V guest VMs can also be manually recovered quite easily in cases where multiple Hyper-V guest VMs share a single LUN in a many-to-one configuration.

The following three examples will demonstrate how to recover a Hyper-V guest when:

- The Hyper-V guest is using a one-to-one guest-to-volume configuration.

- Multiple Hyper-V guests share a common volume (many-to-one configuration).
- A Hyper-V guest VM is running on a cluster shared volume (CSV).

5.1.1 Recover a Hyper-V Guest VM with One-to-One Volume Configuration

The following procedure will demonstrate how to recover a Hyper-V guest VM given the following configuration:

- A standalone Hyper-V host server.
 - A Hyper-V guest VM is configured with one or more virtual hard disks files on a single dedicated SAN volume mapped to the host as LUN 1 (with a drive letter or a mount point).
 - A Dell Compellent Replay schedule is set to take Replays of the volume per a predetermined schedule (e.g. hourly, daily, etc.)
- 1) Just in case the guest VM settings need to be configured from scratch (this should not be necessary, but just in case...) record the details about the guest VM's CPUs, RAM, Virtual networks, IP addresses, etc.
 - 2) Also document the disk letter or mount point information for the volume to be recovered.
 - 3) Power down the guest VM to be recovered if it is not already powered off.
 - 4) If using a mount point, use Disk Manager on the host server to remove the mount point from the volume.
 - 5) Use Disk Manager to take the disk off line.
 - 6) Log on to the Dell Compellent Storage Center Manager GUI where the volume resides.

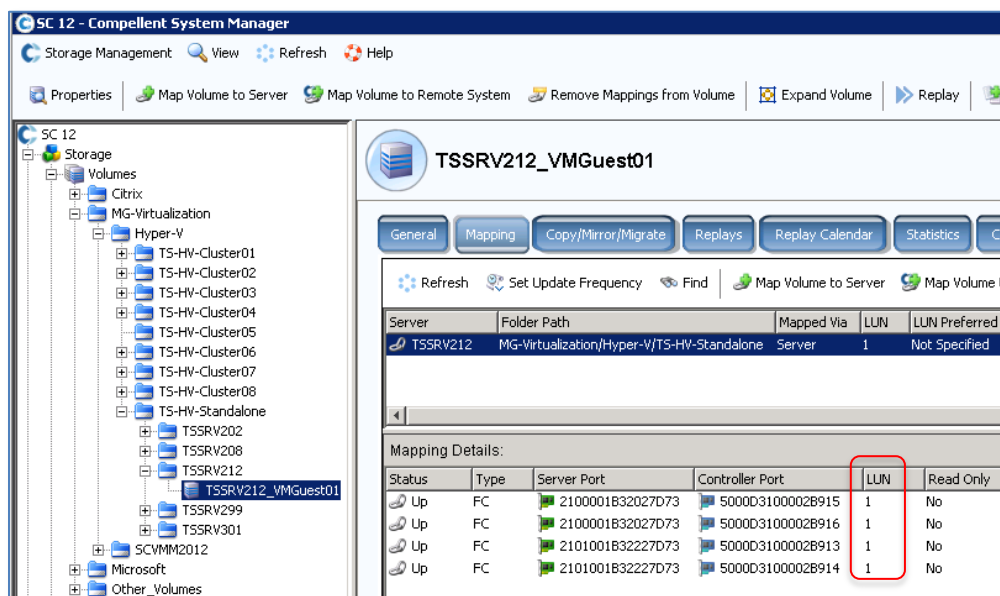


Figure 45. Verify LUN Number

- 7) Select the volume containing the Hyper-V guest VM's virtual hard disk files.
- 8) Select the **Mapping** tab.

9) Verify the mapping details and record the LUN number shown in Figure 45.

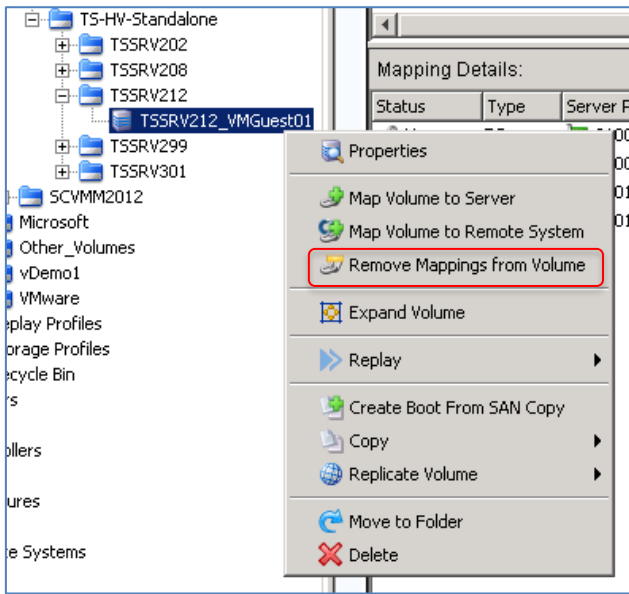


Figure 46. Remove Mappings from Volume

- 10) Right click on the volume in the left pane and select **Remove Mappings from Volume**.
- 11) Once the volume is unmapped, use Disk Manager on the server host to rescan the disks to make sure the host no longer sees it.
- 12) Select the **Replays** tab.

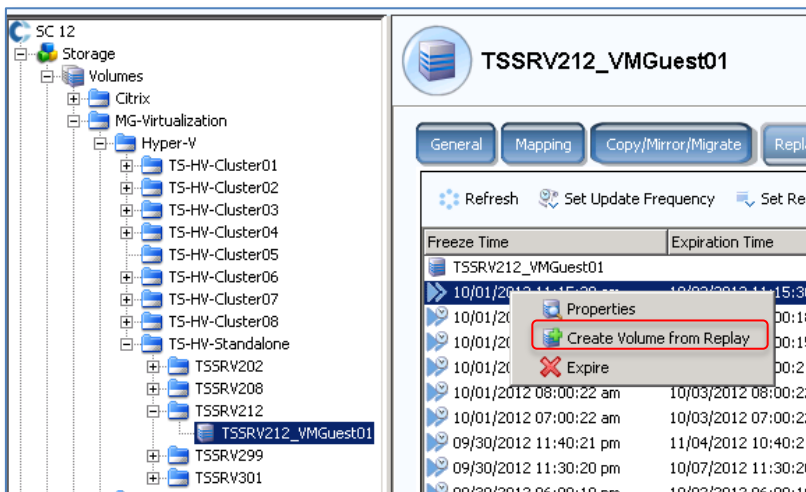


Figure 47. Create Volume from Replay

- 13) Right click the desired Replay and select **Create Volume from Replay**.

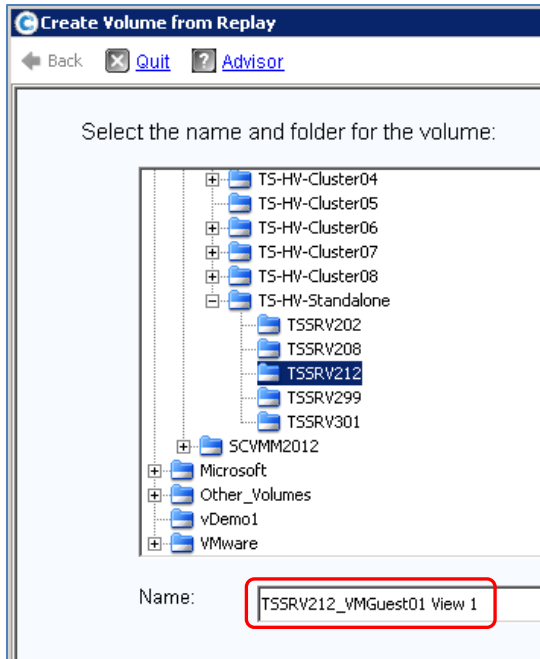


Figure 48. Provide a Name for the View Volume

- 14) Give the Replay View Volume a descriptive name or choose the suggested default name of **<volumename> View 1** and then click the **Create Now** button.

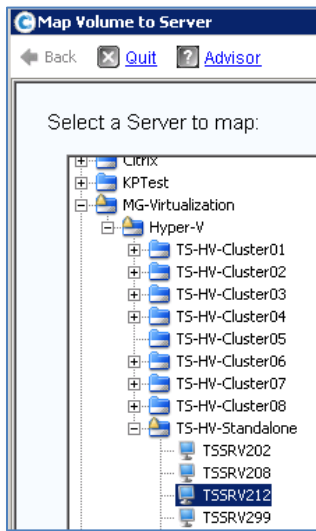


Figure 49. Select the Server to Map the View Volume to

- 15) Select the server to map the View Volume to and click on the **Continue** button. In this example, the View Volume will be mapped back to the original host server.

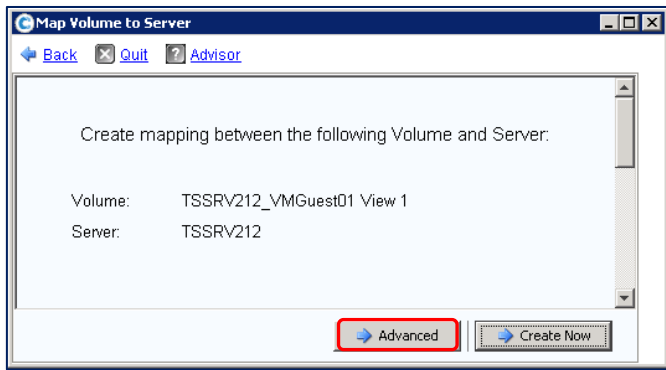


Figure 50. Set Advanced Options for Volume Mapping

16) Click on the **Advanced** button.

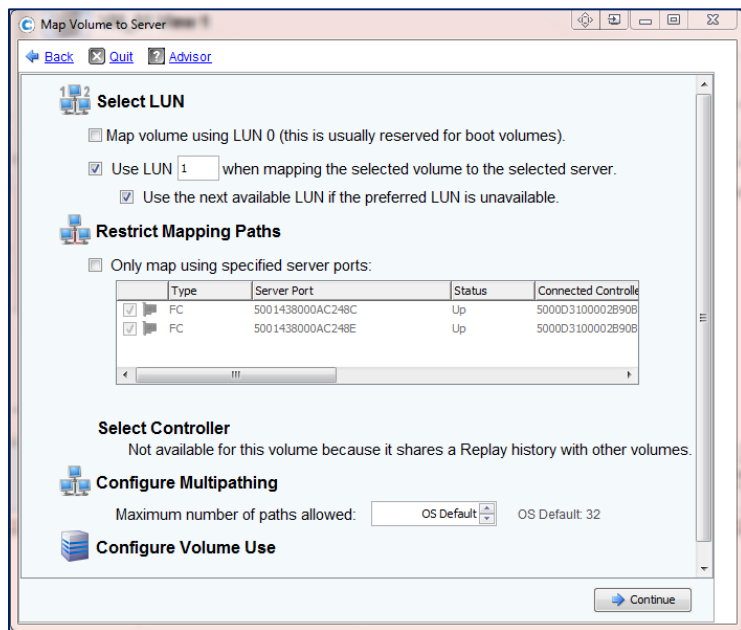


Figure 51. Provide a LUN number for the View Volume

- 17) Check the box for **Use LUN** and provide the same LUN number used by the original volume. Since the LUN is being mapped to the same host, keeping everything about the LUN the same (the LUN number and the drive letter) will allow the guest VM to boot up without any need for reconfiguration of Hyper-V settings on the host server.
- 18) Click on **Continue**, then on **Create Now** to finish mapping the View Volume to the server.
- 19) Configure Data Instant Replay settings for the View Volume by right-clicking on the volume and selecting **Replay**→**Configure Data Instant Replay**.
- 20) If remote replication is being used, configure the View Volume to replicate to another Storage Center as desired.
- 21) Once the View Volume has been mapped to the host server, return to **Disk Manager** on the host server and perform a disk rescan.

- 22) Bring the new disk online, and verify the volume displays properly and that has the same drive letter as the original volume as noted in step 2 above (if not, assign a drive letter or change the drive letter). If using mount points, recreate the mount point so that it uses the same name and location as noted above in step 2.
- 23) Using Hyper-V Manager, start the recovered Hyper-V guest VM and verify functionality.
- 24) Finally, perform any clean-up for the environment as desired. For example, rename the View Volume and remove any unneeded files, mappings or volumes.

5.1.2 Recover a Hyper-V Guest VM with a Many-to-One Volume Mapping

To recover a Hyper-V guest VM from a LUN that this also used by other guest VMs in a many-to-one configuration, the process is similar to the steps in 5.1.1 above. However, because it may not be possible or desirable to take the other guests that are sharing that volume off line, a different approach is necessary to recover a single guest VM.

Instead of removing the original LUN mapping and replacing it with a View Volume of a Replay (which would affect all of the guest VMs), the original LUN is left in place (so the other guest VMs can stay online), and a View Volume is mapped as a new LUN to the host. Once the new LUN is mapped to the host server, there are a couple of different recovery options to recover a single guest VM:

- Copy the guest VM's virtual hard disk files (VHD/VDX) from the new LUN back to the original locations on the original LUN using a Windows file copy operation. Depending on the size of the virtual hard disk files, this may not be practical if they are extremely large.
- Use Hyper-V Manager to either recreate (Server 2008 Hyper-V) or import (Server 2012 Hyper-V) the guest VM on the new volume.

The following example will demonstrate how to perform either of these two operations.

Environment:

- A standalone Hyper-V host server.
- Four hyper-V guest VMs named GuestVM-01 – 04 that are sharing the same LUN for their virtual hard disk and configuration files.
- A Dell Compellent Replay schedule is set to take Replays of this LUN according to a predetermined schedule (e.g. hourly, daily, etc.).
- GuestVM-04 has experienced an event requiring it to be recovered to a previous known-good state.

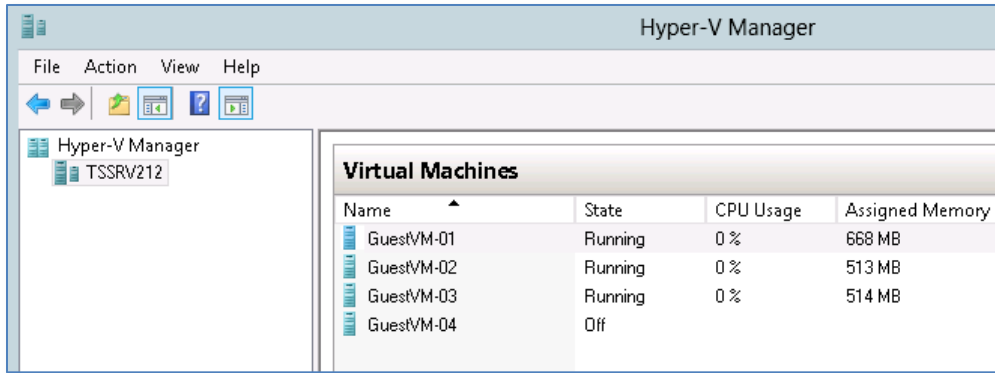


Figure 52. Power Off the Guest VM to be Recovered

- 1) Document the settings for the guest VM so it can be recreated in Hyper-V Manager (e.g. number of CPUs, RAM configuration, Virtual networks, IP addresses, etc.).
- 2) Document the existing location of the guest VM's files on the host server (mount point or drive letter along with paths to the virtual hard disk and configuration files).
- 3) Power off the guest VM to be recovered (in this example, GuestVM-04).
- 4) Log on to the Dell Compellent Storage Center Manager GUI.
- 5) Locate the volume containing the Hyper-V guest VM files to be recovered.

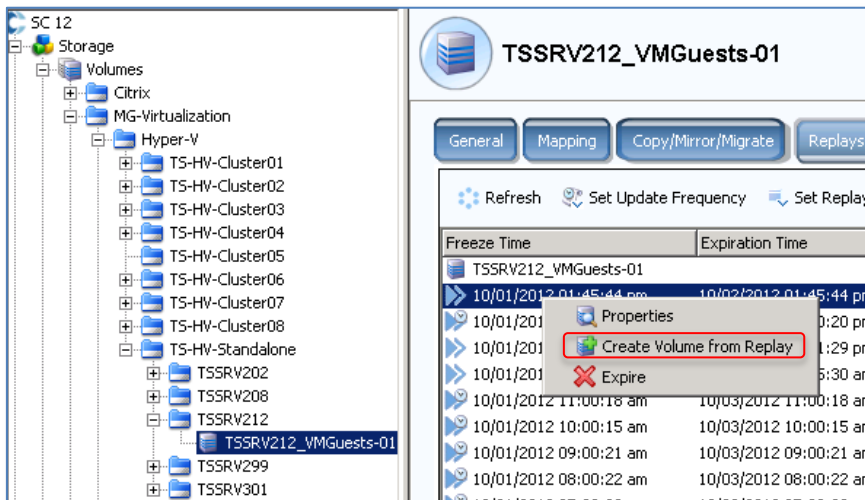


Figure 53. Create Volume from Replay

- 6) Click on the **Replays** tab.
- 7) Right click the desired Replay and select **Create Volume from Replay**.

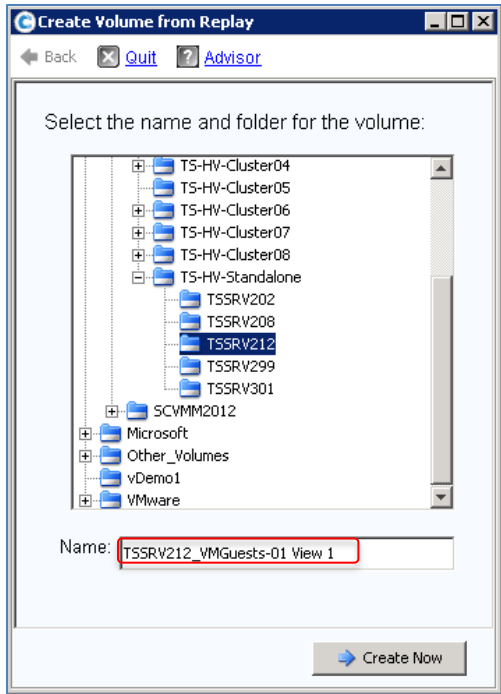


Figure 54. Provide a Name for the View Volume

- 8) Give the View Volume a descriptive name or choose the suggested default name of **<volumename> View 1** and then click the **Create Now** button.

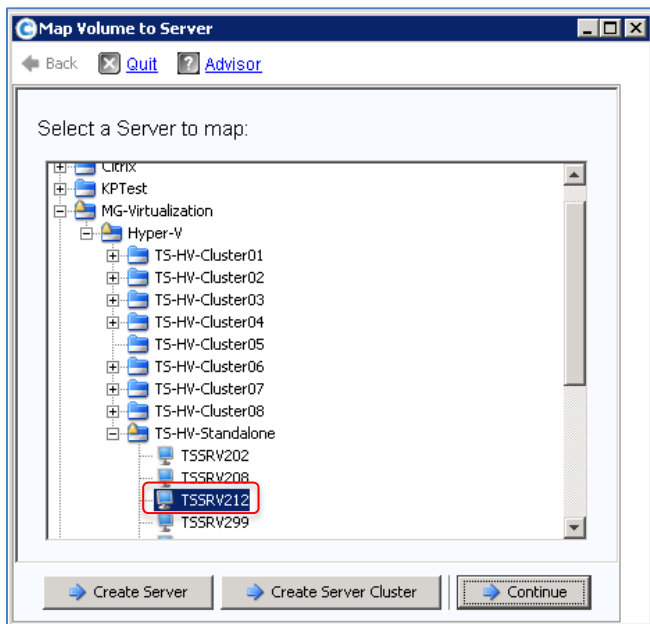


Figure 55. Select the Server to Map the View Volume to

- 9) Select the server to map the view volume to and click on the **Continue** button. This can be the same host server or a different host server. In this example, the View Volume will be mapped back to the original host server.
- 10) Click on **Create Now** to finish mapping the view volume to the server.

Disk 0 Basic 125.00 GB Online	System Reserved 350 MB NTFS Healthy (System, Active, Primary Part	Server2012-DC_VL_Boot (C:) 124.65 GB NTFS Healthy (Boot, Page File, Crash Du
Disk 1 Basic 500.00 GB Online	VMGuests-01 (D:) 500.00 GB NTFS Healthy (Primary Partition)	
Disk 2 Basic 500.00 GB Online	VMGuests-02 (E:) 500.00 GB NTFS Healthy (Primary Partition)	

Figure 56. New View Volume mapped to a Server Host and Renamed

- 11) Once the View Volume is mapped and available on the host, launch **Disk Manager** on the host and rescan the disks.
- 12) The newly mapped View Volume should be listed as a new disk. Bring the disk online, rename it, and assign it a drive letter or mount point. In this example, it is assigned drive letter E, and is renamed as **VMGuests-02**.

Note: If **automount** is enabled, the disk should be assigned a drive letter automatically. If not, manually assign the new View Volume a drive letter or assign it a mount point.

- 13) Chose a recovery option:
 - Option 1: Copy the guest VM's virtual hard disk (VHD/VHDX) file from the new volume (the E drive in this example) back to its original location on the host server (see step 14 below). While copying a virtual hard disk file back to its original location is usually the preferred option, if it is extremely large it may not be practical.
 - Option 2: Use Hyper-V manager to recreate (Server 2008 R2) or import (Server 2012 Hyper-V) the guest from its new location on the new volume (see step 15 below).

Option 1:

- 14) Decide what to do with the original virtual hard disk file (rename it, move it, or delete it). If renaming the file, verify that the volume has enough space for the old and the new virtual hard disk.
 - a. Copy the recovery virtual hard disk file back to its original location using Windows Explorer (in this example, from the E drive to the D drive)
 - b. After the virtual hard disk has finished copying, rename the virtual hard disk file so that it is different than original file name. For example:

Original name: "GuestVM-04_W2K8R2_boot.vhdx"
 New Name: "GuestVM-04_W2K8R2_boot_**new**.vhdx"

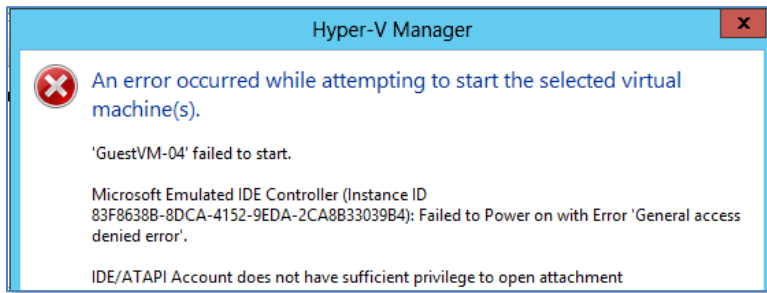


Figure 57. Rename the Recovered VHD File to Avoid Access Denied Errors

Note: Renaming the virtual hard disk and re-associating it with the guest VM with Hyper-V Manager is necessary to allow the guest VM to start without permissions errors.

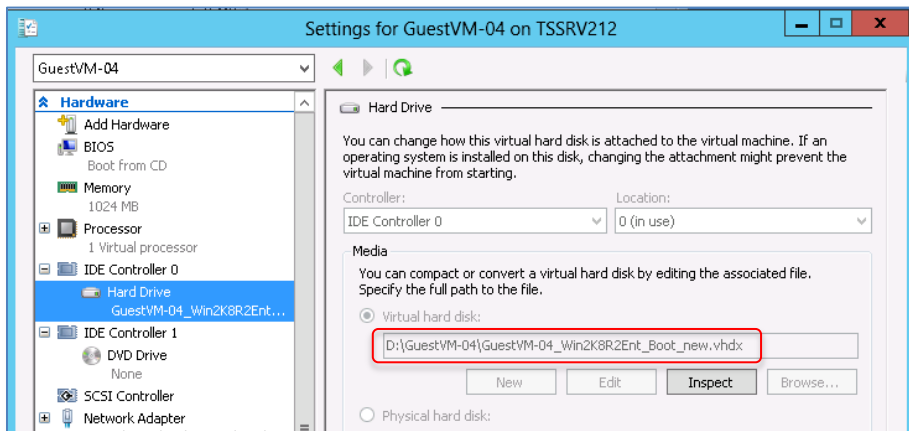


Figure 58. Use Hyper-V Manager to Repoint to the Recovered VHD file

- c. Use Hyper-V Manager to repoint the guest VM to the renamed recovery virtual hard disk file. No other changes in Hyper-V Manager should be necessary.
- d. Start the guest VM and verify functionality.
- e. Perform any clean-up steps on the SAN and the host server.

Option 2:

- 15) Option 2: Use Hyper-V Manager to Import the guest VM's settings from the new volume

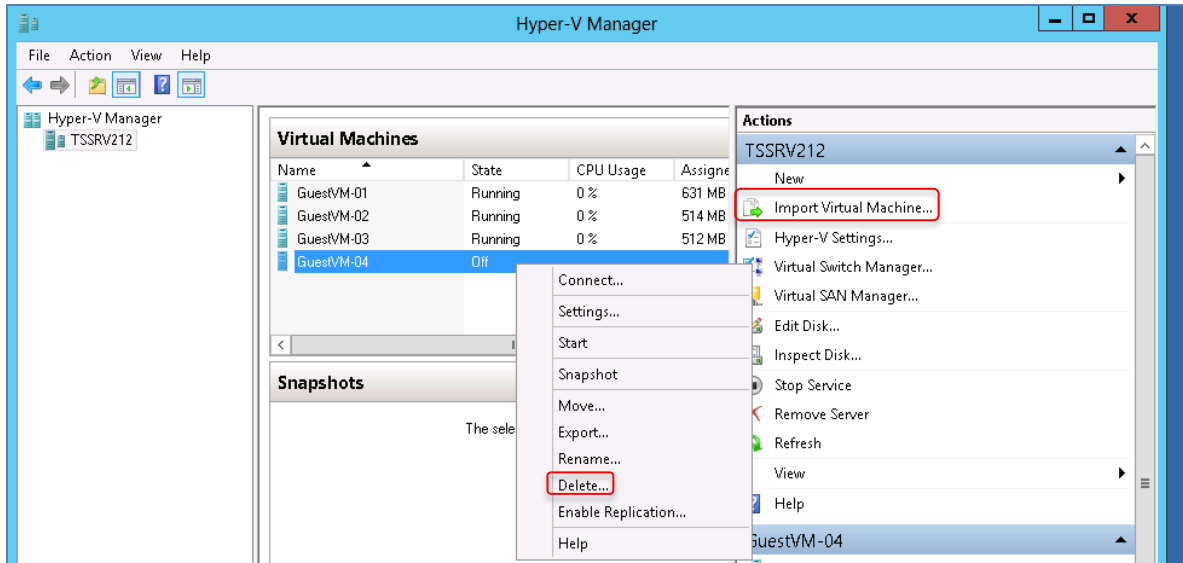


Figure 59. Delete the Old Guest, Then Run the Import Virtual Machine Wizard

- a. Use Hyper-V Manager to delete the guest VM to be recovered.
- b. If running Server 2008 R2 Hyper-V, recreate the guest from scratch by creating a new virtual machine, pointing to the guest’s virtual hard disk file on the new volume. When finished, start the guest, and configure any necessary IP addresses for the guest, and perform any other clean-up needed on the host or the SAN.
- c. If running Server 2012 Hyper V, use Hyper-V Manager to import the recovered guest VM. Click on **Import Virtual Machine** and click on **Next**.

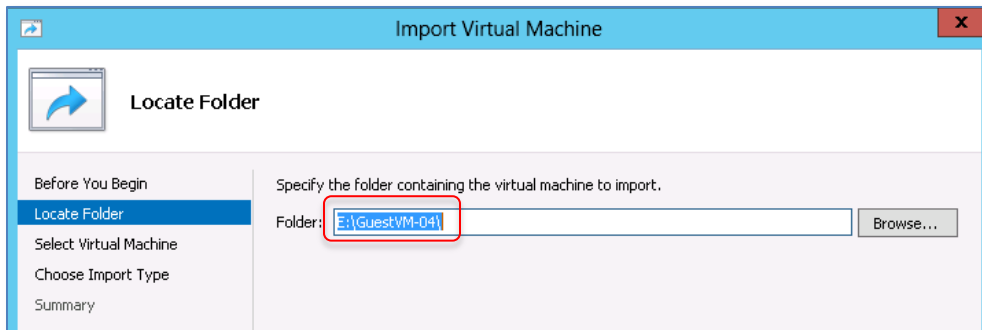


Figure 60. Specify the Folder Containing the Guest VM to import

- d. Select the folder containing the guest VM on the new volume, in this example, E:\GuestVM-04. Then click on **Next**.

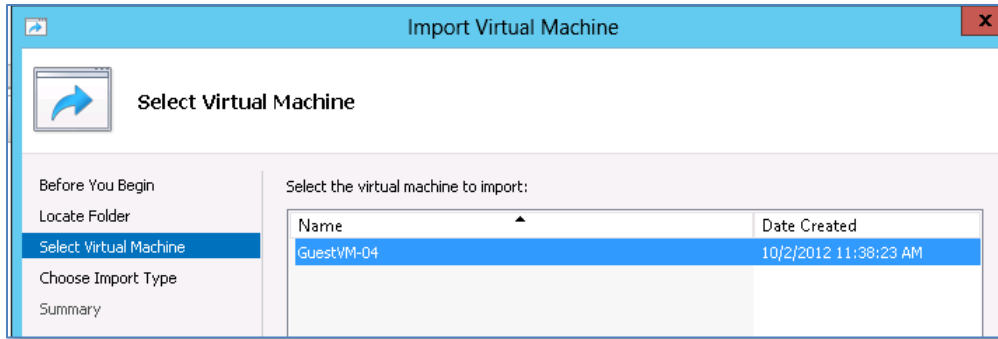


Figure 61. Verify the Name of the Guest VM to Import

- e. Confirm the name of the guest VM to be imported, and then click on **Next**.

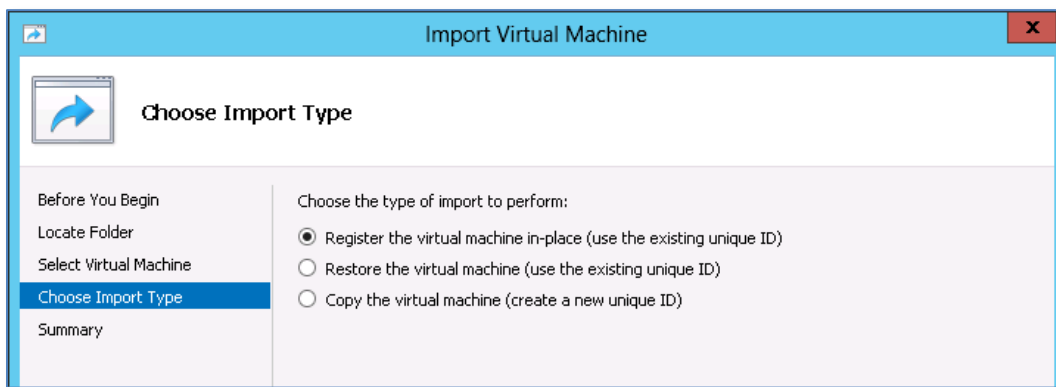


Figure 62. Choose the Import Type

- f. Select the **Import Type** as **Register the virtual machine in-place** and then click on **Next**.

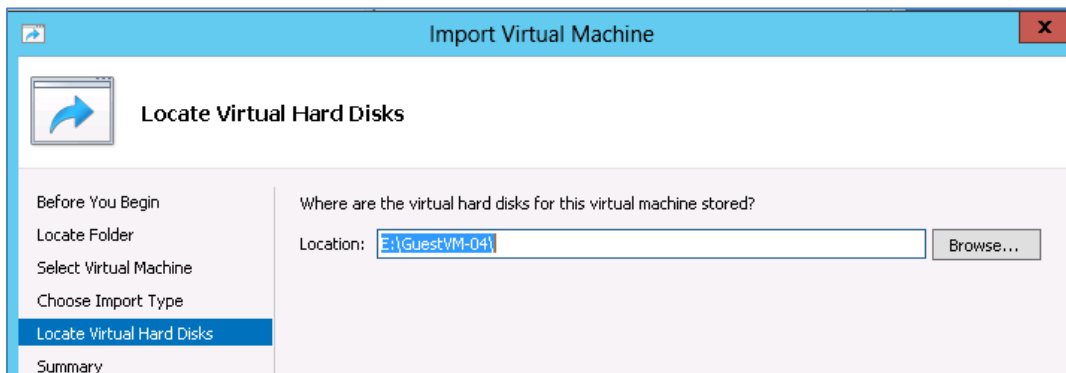


Figure 63. Specify the Location of the Guest VM's Virtual Hard Disk Files

- g. For **Locate Virtual Hard Disks**, specify the path to the virtual hard disk on the new volume. Then click on **Next**. In this example, the path is E:\GuestVM-04\.
- h. On the summary screen click on **Finish**.
- i. Right-click on the guest VM in Hyper-V Manager and verify the guest VM's settings for memory, CPUs, and virtual hard disks.

- j. Start the guest VM and verify functionality.
- k. Clean up the environment as required to remove any unneeded files, mappings or volumes.

5.1.3 Recover a Hyper-V Guest VM that resides on a Cluster Shared Volume

The process of using Storage Center Replays to recover guest VMs that reside on Cluster Shared Volumes (CSVs) is similar to the process of recovering guest VMs to standalone hosts, as detailed in the preceding sections 5.1.1 and 5.1.2. However, there is one additional step required when dealing with CSVs - changing the disk signature.

Windows servers assign each volume a unique disk ID (or signature). For example, the disk ID for an MBR disk is an 8-character hexadecimal number such as 045C3E2F4. No two volumes mapped to a server can have the same disk ID.

When a Storage Center Replay is taken of a Windows volume, the Replay is an exact point-in-time copy of the volume, which includes the Windows disk ID information. Therefore, if a View Volume is created from that Replay and presented back to the same server or cluster, it will cause a disk ID conflict.

With standalone Windows servers, disk ID conflicts are avoided because standalone servers are able to automatically and dynamically detect duplicate disk IDs and change them. No user intervention is required.

However, Windows is not able dynamically change conflicting disk IDs when disks are configured as CSVs.

Note: this issue exists with both Server 2008 R2 Hyper-V and Server 2012 Hyper-V.

When attempting to map a View Volume of a CSV back to any server in the same cluster, the new volume (due to having a duplicated disk ID) conflicts with the existing CSV, and the new volume stays stuck in a reserved and unmanageable state.

Note: Duplicate disk IDs on a cluster may cause cluster instability resulting in a service outage, data loss or corruption.

There are a couple of ways of working around the duplicate disk ID issue as detailed below.

Option 1 – Map the View Volume of the CSV to another host that is outside of the cluster, and copy the guest VM's files over the LAN to recover the guest (see section 5.1.1 above).

Option 2 – Map the View Volume to another Windows host outside of the cluster and use Diskpart.exe to change the disk ID, and then re-map the View Volume to the Cluster. The

steps to use Diskpart to change the disk ID are detailed in Section 5.1.4 below.

5.1.4 Procedure for Assigning a New Disk ID to a View Volume

- 1) Logon to the standalone Windows server host that the View Volume of the desired CSV will be mapped to.
- 2) Open a command window and then type **Diskpart.exe** and press **Enter**.

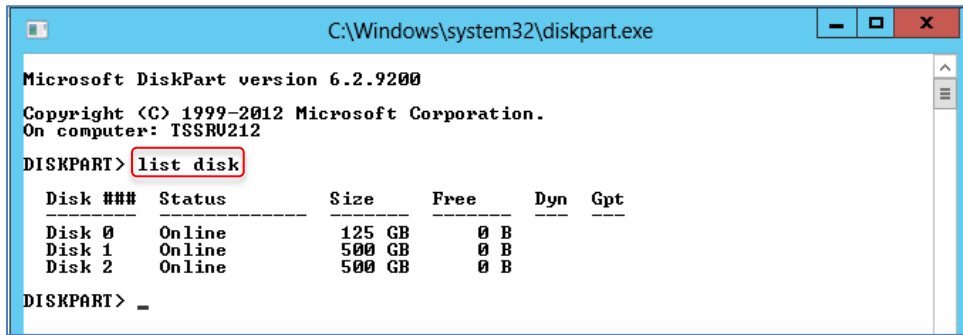


Figure 64. Diskpart Command Window

- 3) Type **List Disk** and press **Enter**.
- 4) Make note of the current list of disks (Disk 0, Disk 1, Disk 2).
- 5) Using Storage Center Manager, map the desired View Volume to this host.
- 6) From the **Diskpart** command prompt, type **rescan** and press **Enter**.




Figure 65. Use Disk Manager to Bring a Disk Online

- 7) Launch **Disk Management** on the host server and bring the disk online.

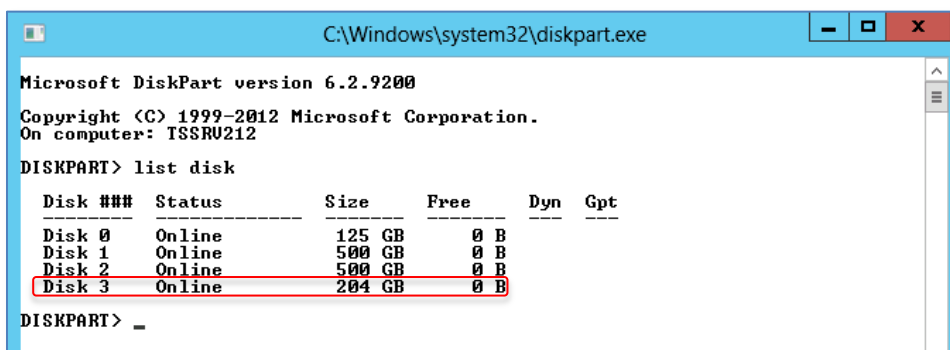


Figure 66. New Volume Now Listed with Diskpart

- 8) Return to the Diskpart command prompt window and type **List Disk** and press **Enter**.
- 9) The new disk (Disk 3 in this example) should now be listed. Usually (but not always) the bottom disk on the list will be the one just added).

```

C:\Windows\system32\diskpart.exe
Microsoft DiskPart version 6.2.9200
Copyright (C) 1999-2012 Microsoft Corporation.
On computer: TSSRU212

DISKPART> list disk

   Disk ###  Status              Size          Free           Dyn  Gpt
   -----  -
   Disk 0    Online              125 GB         0 B
   Disk 1    Online              500 GB         0 B
   Disk 2    Online              500 GB         0 B
   Disk 3    Online              204 GB         0 B

DISKPART> select disk 3
Disk 3 is now the selected disk.
DISKPART> uniqueid disk
Disk ID: 12341256
DISKPART> uniqueid disk ID=1234abcd
DISKPART> uniqueid disk
Disk ID: 1234ABCD
DISKPART> _

```

Figure 67. Diskpart Select Disk Command

- 10) Referring to Figure 67, type **select disk #** (where **#** represents the number of the new disk, in this example, disk 3) and then press **Enter**.
- 11) Type **uniqueid disk** and press **Enter** to view the current unique ID for the disk. In this example, the unique ID is **12341256** as shown in Figure 67.
- 12) To change the disk ID, type **uniqueid disk ID=<newid>** and press **Enter**.
 - For **<newid>** provide a new random ID of your choice. For an MBR disk, the new ID must be an 8-character string in hexadecimal format using a mix of the numbers 0 – 9 and the letters A – F.
 - For a GPT disk, the new ID must be a GUID (Globally Unique Identifier).
- 13) Type **uniqueid disk** again and press **Enter** to verify the ID is now changed.
- 14) Now that View Volume has a new signature, it can now be unmapped from the standalone host server and re-mapped to the cluster without causing a disk ID conflict.
- 15) Once it has been mapped to the cluster, the guest VM can be manually recovered as demonstrated in Sections 5.1.1 and 5.1.2 above.

5.2 Use Storage Center Replays to Create a Test Environment

Storage Center Replays can be used to create space efficient test and development environments of production Hyper-V guest VMs by mapping View Volumes of Replays to other (usually isolated) server hosts or clusters. These hosts and clusters can be local (in the same data center) or remote (when SAN Volumes and their Replays are replicated to another Storage Center at another location).

Note: To avoid IP, MAC or server name conflicts, a separate network should be created to isolate cloned servers in a test or development environment.

The procedure to create a test environment with cloned Hyper-V guests from View Volumes of Dell Compellent Replays is very similar to the guest VM recovery process previously shown in sections 5.5.1 and 5.5.2. But instead of replacing an existing guest VM, simply create a new (or import) a guest VM to the desired server host or Hyper-V cluster.

5.3 Use Storage Center Replays to Create New Hyper-V Guests from a Gold Image

With Dell Compellent Storage Center Replays, an administrator can use View Volumes to quickly deploy new Hyper-V guest VMs from a gold image as shown in Figure 68 below. Deploying new Hyper-V guest VMs from gold images represents a “best of all worlds” design because:

- New guest VMs can be deployed from a gold image very quickly with minimal reconfiguration.
- SAN usage is maximized - if the four guest VMs were deployed below in Figure 68 without using a gold image, they would each consume a full 16 GB space for a total of 64 GB. By using a gold image, the same four guest VMs consume only 20 GB initially (the source volume plus 4 new View Volumes that represent only the changed data for each guest).
- New writes go to Tier 1-RAID 10 for maximum write performance.
- Data files on the source gold image that don't change can take advantage of the read performance of RAID 5, and can also reside on lower tier disk for additional SAN efficiency and cost savings.
- All data is subject to Data Progression which will move data up or down between tiers and RAID levels as needed to ensure an optimal balance between I/O performance and SAN efficiency.
- Will Dell Compellent, this all happens automatically without any user intervention.

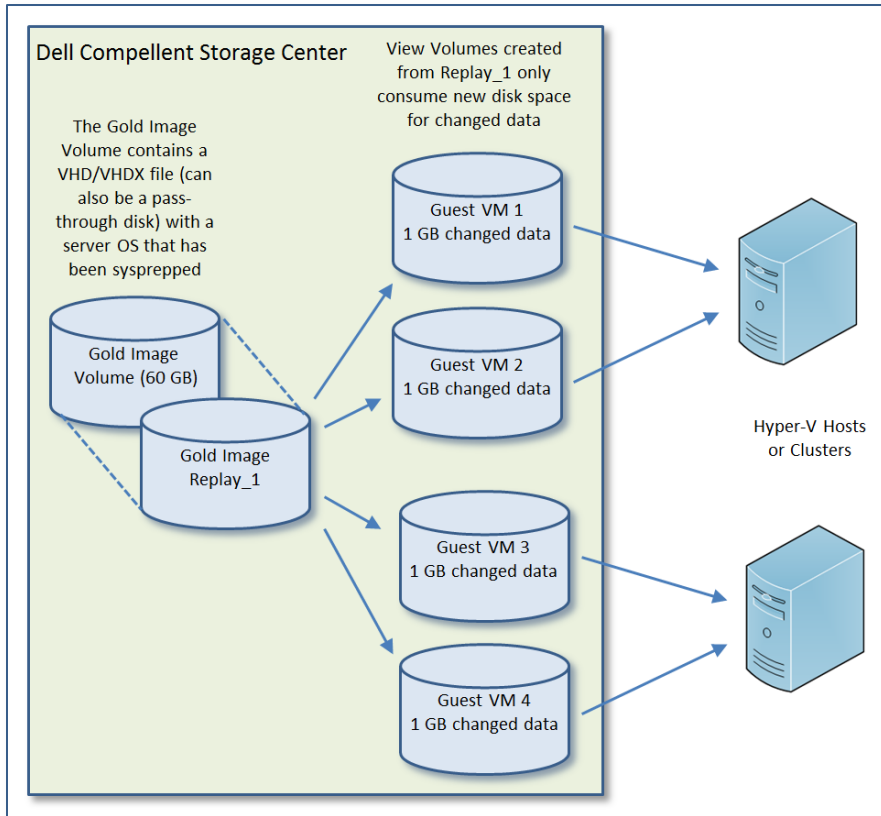


Figure 68. Gold Image Used for View Volumes for Multiple Hyper-V Guests

In the following example, a VHD/VHDX file is used as the boot volume (but pass-through boot volumes will also work). It also assumes that the guest VM is running a Windows Server OS that is version 2003 or newer.

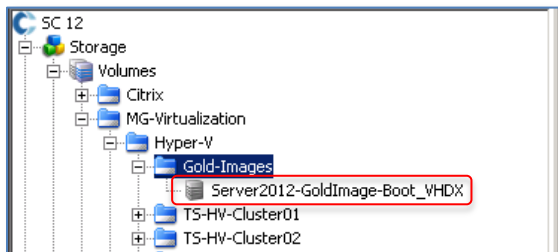


Figure 69. Create a Gold Image Volume

- 1) As shown in Figure 69, from the Dell Compellent Storage Center Manager GUI, create a new volume with the appropriate size for the gold image in a location of your choice. In this example, a **Gold-Images** folder was created as a place to hold this and future gold image volumes.
 - For server 2003: create at least a 35 GB volume to hold a 30 GB virtual hard disk.
 - For Server 2008/R2: create at least a 45GB volume to hold a 40 GB virtual hard disk.
 - For Server 2012: create at least a 65 GB volume to hold a 60 GB virtual hard

disk.

- Factors such as the amount of RAM to be assigned to the guest VMs and the location of the page file may require a larger gold image virtual hard disk size.
- If Hyper-V based snapshots will be taken of these guests, create a larger volume to accommodate the differencing virtual hard disks that Hyper-V based snapshots create.
- Since Dell Compellent uses thin provisioning, it is recommended to start with a large enough gold image virtual hard disk and volume so that they will not have to be expanded later.




Figure 70. Format the Volume for the Gold Image

- 2) Don't specify a Replay schedule for the volume at this time (a manual Replay will be taken later).
- 3) Map the volume to a to a Hyper-V host server and use Disk Manager to bring the disk on line, format it, assign it a label, and assign it a mount point or a drive letter. In this example, it is assigned as the G drive. The volume (65 GB) will be the target location for a new Server 2012 guest that will be configured to use a 60 GB (dynamically expanding) VHDX file.

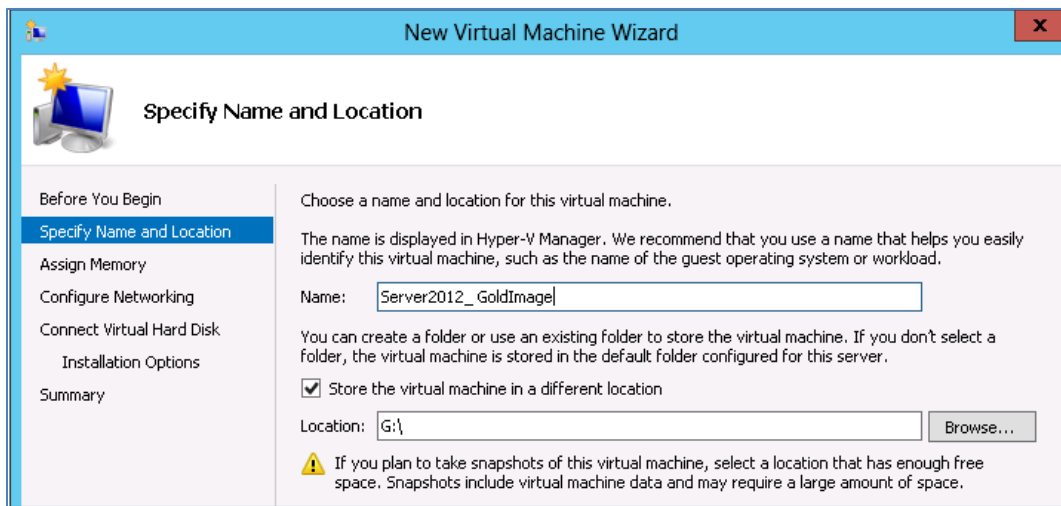


Figure 71. New Virtual Machine Wizard

- 4) Using Hyper-V Manager, launch the **New Virtual Machine Wizard** and create a new guest VM and point its configuration and VHD/VHDX location to the root of the new volume, in this example, the G drive.

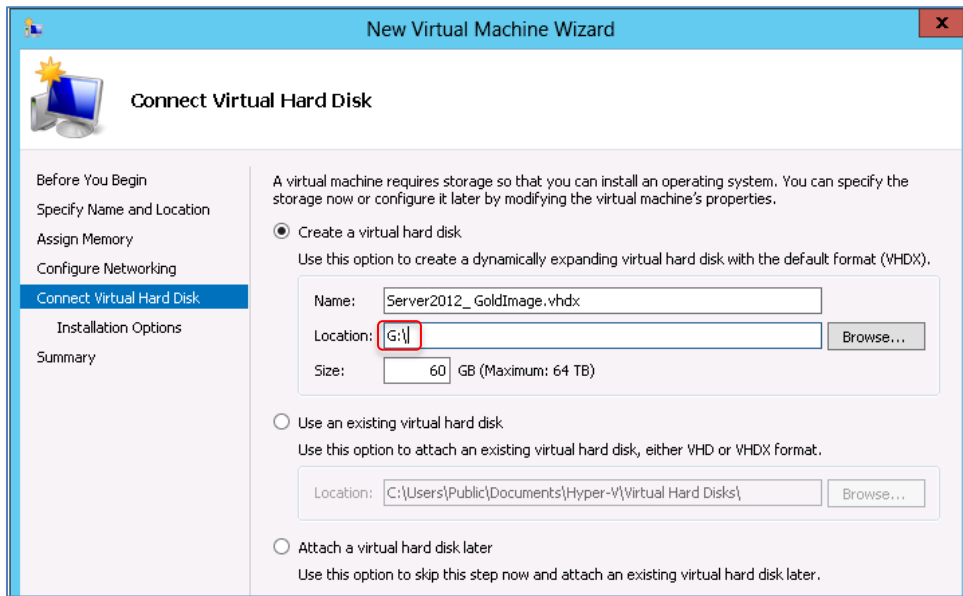


Figure 72. Create a New Virtual Hard Disk

- 5) As part of creating the new virtual guest, create a new virtual hard disk on this volume of the type and size desired. In this example, the Server 2012 Hyper-V wizard defaults to a dynamically expanding VHDX file format.
 - a. Provide an intuitive name that will help identify the purpose of the virtual hard disk file later.
 - b. Set the path to the root of the volume (the G drive in this example)
 - c. Set the desired size (60 GB in this example).

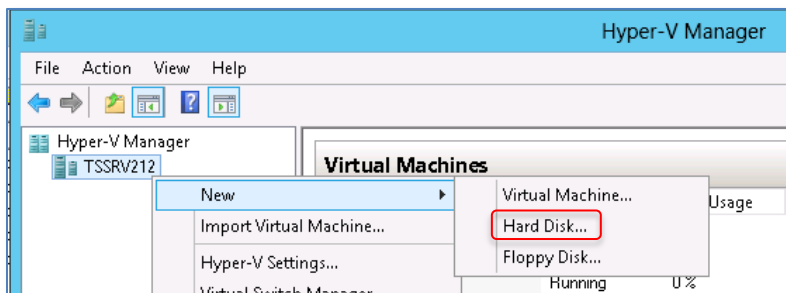


Figure 73. Use Hyper-V Manager to Create a New Virtual Hard Disk

Note: If a different kind of virtual hard disk is desired (for example, a fixed VHD instead of a VHDX), select **Attach a virtual hard disk later**. After completing the **New Virtual Machine Wizard**, use Hyper-V Manager to create a new virtual hard disk file at the root of the desired drive (the G drive in this example). Then edit the guest VM's setting and under disk settings, attach to the virtual hard disk just created.

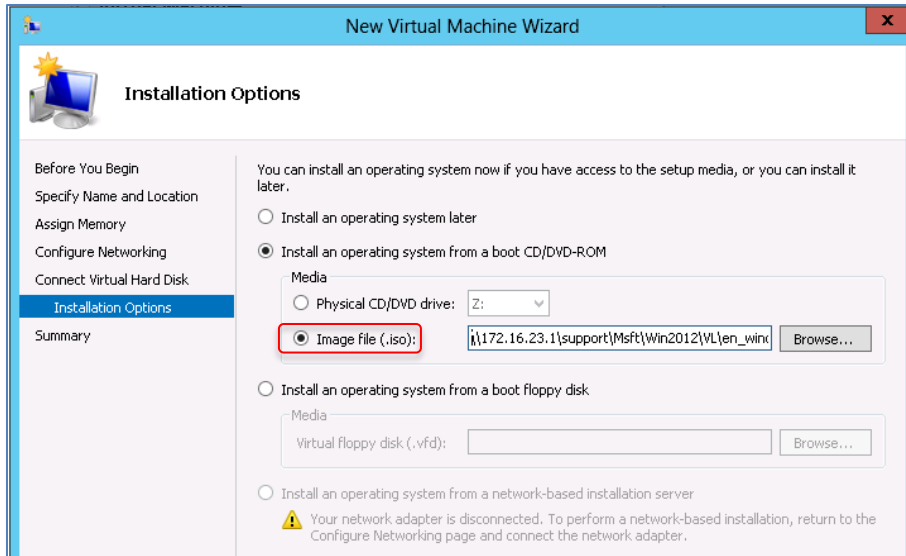


Figure 74. Specify the Location of the Setup Media

- 6) Specify the location of your setup media. In this example, a bootable ISO file is selected as the source media.
- 7) Complete the **New Virtual Guest VM** wizard.

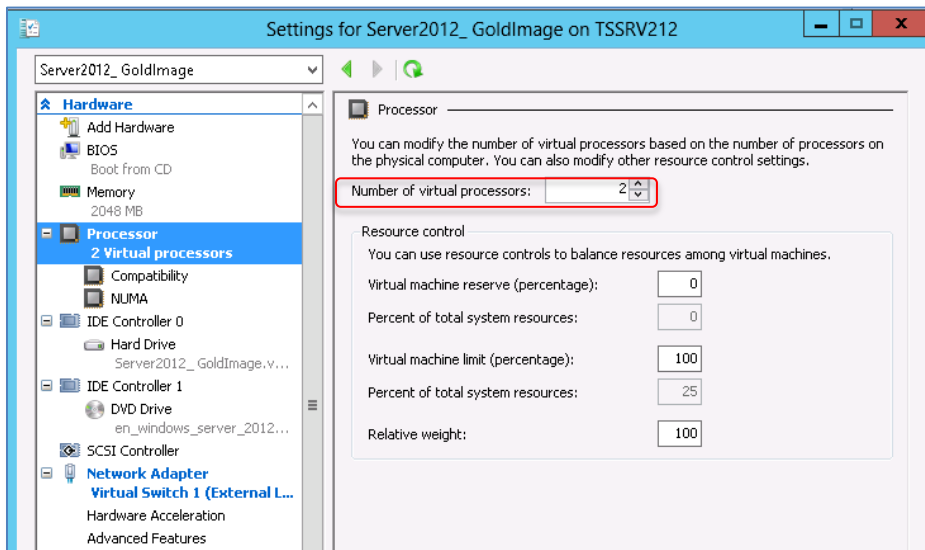


Figure 75. Configure the Guest VM Settings

- 8) Right click on the new guest VM and select **Settings**.
 - a. Configure BIOS and memory settings.
 - b. For processors, select at least 2 virtual processors so that the guest VM will use the multiprocessor Hardware Abstraction Layer (HAL). This will allow both single and multiprocessor Hyper-V guest VMs to be created from this same gold image.
 - c. Verify the virtual IDE controller and disk settings, and select a virtual network adapter (if desired).

- 9) Start the guest VM, boot to the installation media, and install the OS. Patch the OS to the desired level.

Note: If desired, an unattended installation file can be created by using the Microsoft Deployment Toolkit (MDT) and Windows Automated Installation Kit.

- 10) Install any desired roles, features, and applications.
- 11) Make any desired customizations or other changes that will be needed by the guests deployed from this gold image. This will minimize the need to make redundant configuration changes later to the guest VMs created from this gold image.

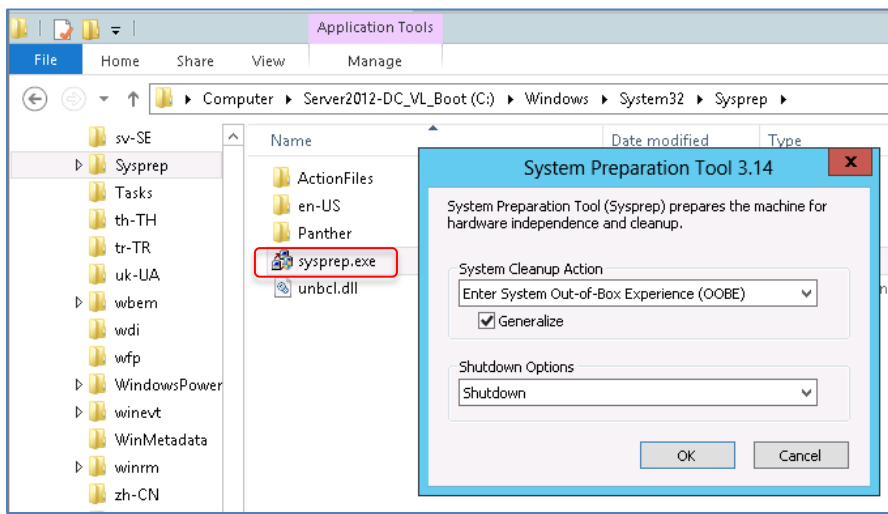


Figure 76. Run Sysprep and Shut Down the Guest

- 12) When the guest VM is built and patched to the desired level run **Sysprep.exe**.
 - a. For **System Cleanup Action**, select **OOBE** with the **Generalize** option checked.
 - b. For **Shutdown Options** select **Shutdown**.
- 13) In Hyper-V Manager, right click on the guest VM and select **Delete**. This will delete the guest VM's configuration files but will preserve the actual boot VHD/VHDX file itself. The VHD/VHDX file is the only file needed on the gold image volume.

Caution: If using Microsoft System Center Virtual Machine Manager (SCVMM), do not delete the guest VM using SCVMM as it will also delete the virtual hard disk file. To prevent the virtual hard disk file from being deleted, first remove the Storage Center mapping to the volume with the gold image VHD/VHDX, and then delete the guest VM from SCVMM.

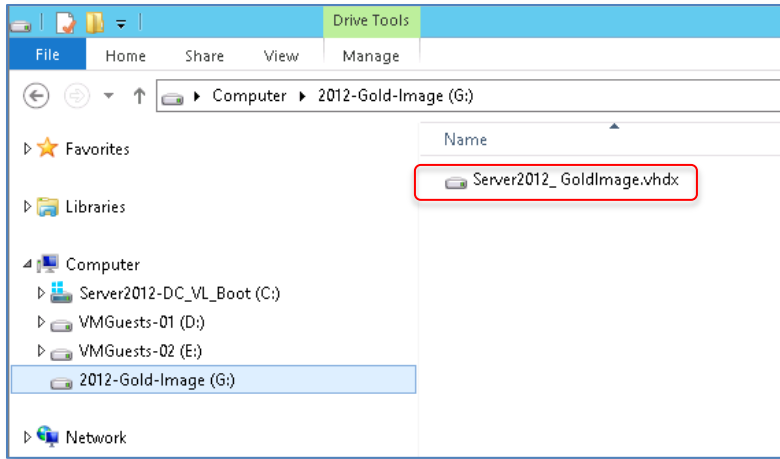


Figure 77. Delete All But the Virtual Hard Disk File from the Gold Image Volume

- 14) Delete all other content on the gold image volume except for the VHD/VHDX file. In this example, the G drive now contains only one file: the gold image VHDX file.
- 15) From the Storage Center Manager GUI, remove the mapping for this volume from the Hyper-V host server.
- 16) Right-click on the volume and select **Replay**, then **Create Replay**.

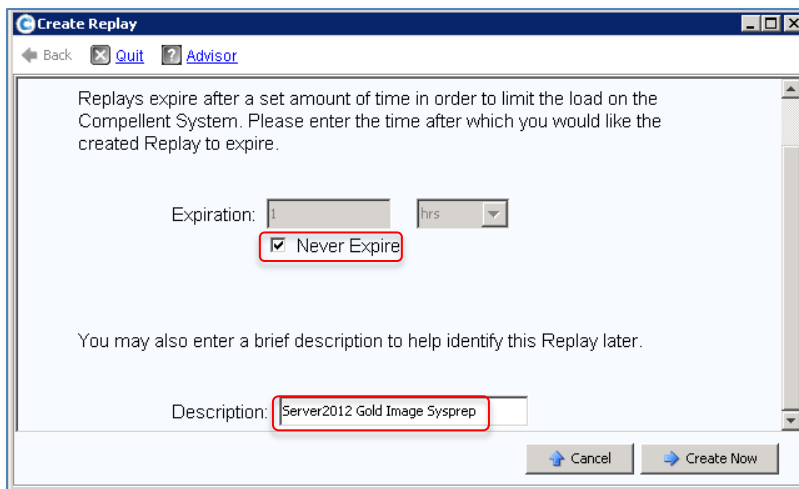


Figure 78. Create Replay of a Gold Image

- 17) Provide a **Description** that identifies the Replay as a gold image and then check the box for **Never Expire**.

18) At this point, View Volumes can now be created from this manual Replay, and mapped to Hyper-V hosts.

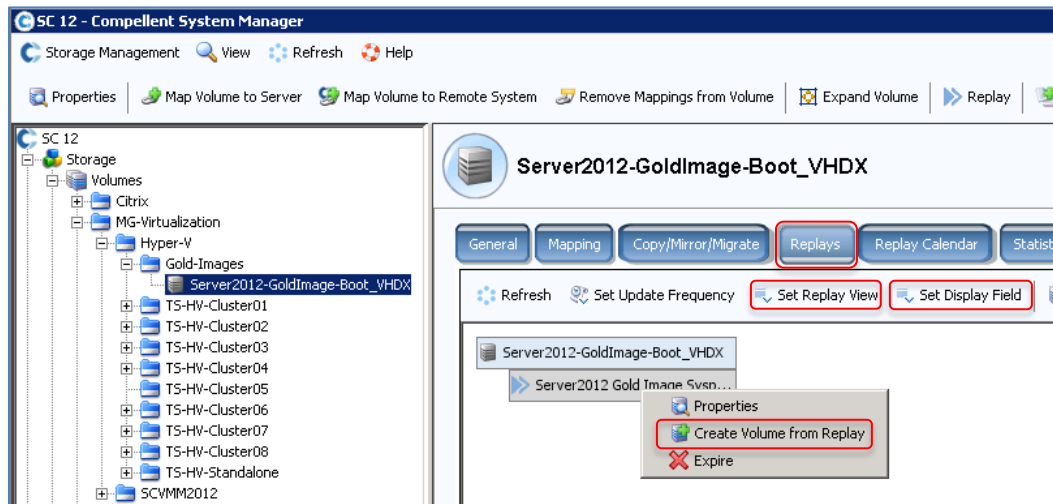


Figure 79. Create a View Volume from a Gold Image

19) Click on the Replays tab

- a. Click on the **Set Replay View** drop-down and change the view to **Show View Volume Tree**.
- b. Click on the **Set Display Field** drop-down and change the Display to **Description**.

20) The manual Replay, which will be the source for View Volumes, will now be displayed right below the source SAN volume it was created from.

- a. Right-click on this manual Replay and select **Create Volume from Replay**.
- b. Provide a name for the volume and provide a volume folder location.
- c. Select a Hyper-V host server of your choice to map the View Volume to and complete the wizard.

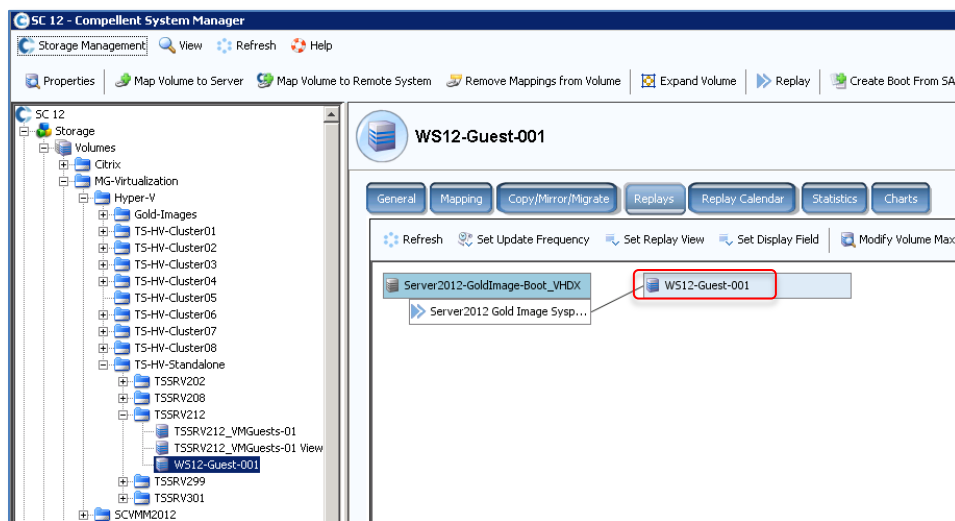


Figure 80. View Volume of a Gold Image Mapped to a Hyper-V Host Server

21) The new View Volume will be displayed (in tree view) as shown in Figure 80.

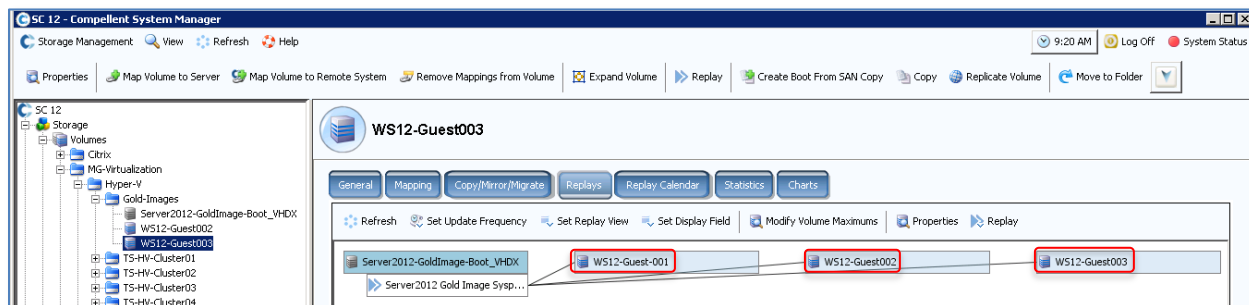


Figure 81. Multiple View Volumes Created from a Single Replay

22) Repeat step 20 above to create additional View Volumes (this can be done now or later). Figure 81 shows how three View Volumes have been created and mapped to a Hyper-V host server. Each View Volume can be mapped to the same or different host servers.




Figure 82. Assign a Drive Letter to the New Guest VM Volume

- 23) Log on to the Hyper-V host Server and use Disk Manager to bring the disk with the syspreped virtual hard disk online, and assign it a drive letter or a mount point. Rename the volume if desired. In this example, the volume was renamed to reflect the name of the new guest VM (WS12-Guest001) and assigned the drive letter P.
- 24) Using Hyper-V Manager, launch the **New Virtual Machine Wizard** to create a new Hyper-V guest VM.

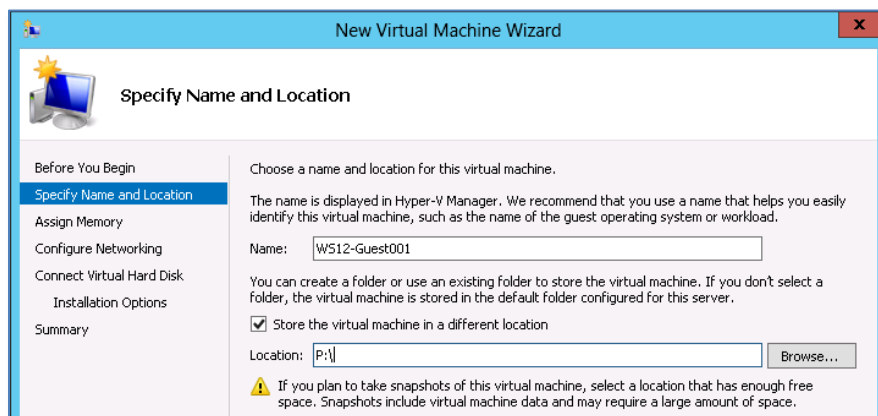


Figure 83. Provide a Name and Location for the Guest VM.

- 25) Specify a name for the new guest VM (in this example, WS12-Guest001).
- 26) Check the box for **Store the Virtual machine in a different location**, click on the

Browse button and select the storage location (the P drive in this example).

27) Specify the memory settings.

28) Select a virtual network adapter (if desired).

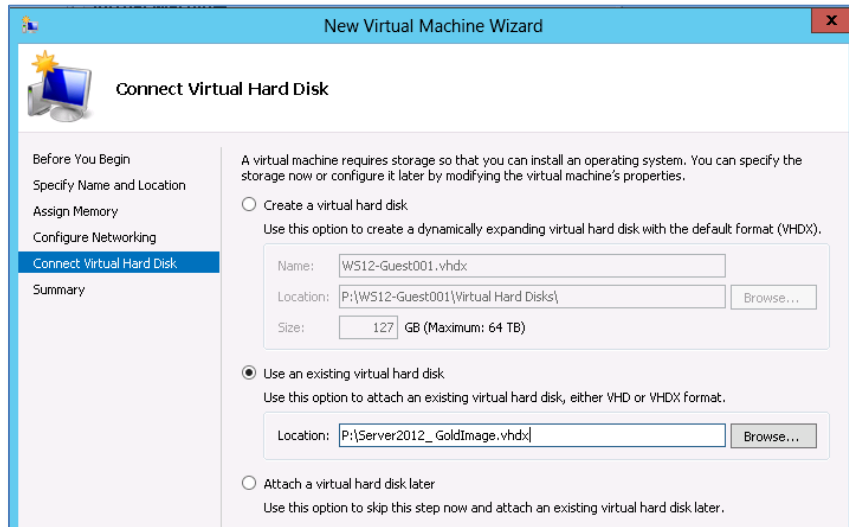


Figure 84. Specify the Location of the Sysprepped Virtual Hard Disk

29) For **Connect Virtual Hard Disk**, click on **Use an existing virtual hard disk**, click on the **Browse** button, and select the sysprepped virtual hard disk file (in this example, located on the P drive).

Note: Rename the virtual hard disk file first if desired.

30) Complete the **New Virtual Machine Wizard**.

31) Using Hyper-V Manager, right-click on the new guest and go to **Settings**.

32) Configure the BIOS, Memory, Processor, Disk and Network Adapter settings as desired.

33) Boot the guest VM and allow the first-time-boot process to complete. Log on to the new Guest VM and configure the guest VM as desired.

34) Repeat steps 23 - 33 for additional View Volumes created from the gold image that are mapped to Hyper-V hosts.

35) Set a Replay schedule of your choice on the View Volumes containing the new guest VMs, and enable replication (if desired) to a remote storage center for both the gold image volume and any View Volumes created from it.

5.3.1 Gold Images and Preserving Balanced Controllers

When using Enterprise Manager or Storage Center Manager to create new volumes on a Storage Center with a dual-controllers, controller ownership of the new volumes is kept in balance automatically. By making sure that each controller "owns" approximately the same number of volumes, it helps ensure an even load on both controllers and their I/O paths.

When one controller in a pair needs to go off line (e.g. for maintenance), volume ownership for all volumes is assumed by the other controller in the pair that is still on line. When both controllers are back on line again, the controllers are “rebalanced” to restore controller volume ownership to a balanced state.

However, when View Volumes are created from a gold image, the View Volumes will always be associated with (owned by) the same controller that owns the gold image volume. Therefore if a large number of View Volumes are created from a gold image, this can result in unbalanced controllers.

To avoid unbalanced controllers when deploying a large number of View Volumes from a gold image, do the following:

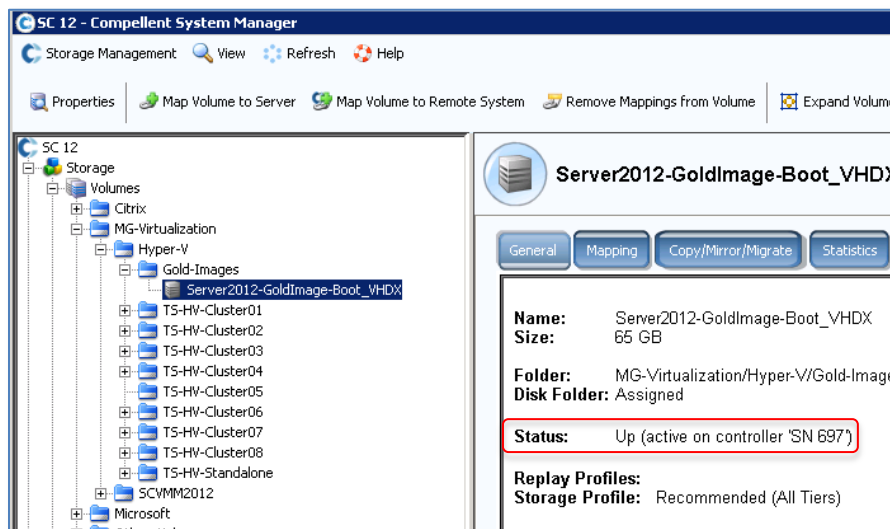


Figure 85. Verify Controller Ownership for a Volume

- 1) Verify the controller ownership for the first gold image volume (created in the previous section 5.3). To verify the ownership, click on the gold image volume in Storage Center Manager and view the **Status** line under the **General** tab. In this example, the gold image created in the previous section was assigned to Controller 697 (in controller pair comprised of 697 and 698).




Figure 86. Add the Controller Number to a Gold Image Volume Name

- 2) Rename the gold image volume to help clearly identify its ownership by Controller 697. In this example, 697 was added to the volume name as shown in Figure 86.
- 3) Create a 2nd gold image volume that can be associated with the other Storage Center

controller (698) in the pair by following the steps shown previously in section 5.3. In this example, the 2nd gold image volume was named similar to the first one, but with 698 added to the name instead of 697.

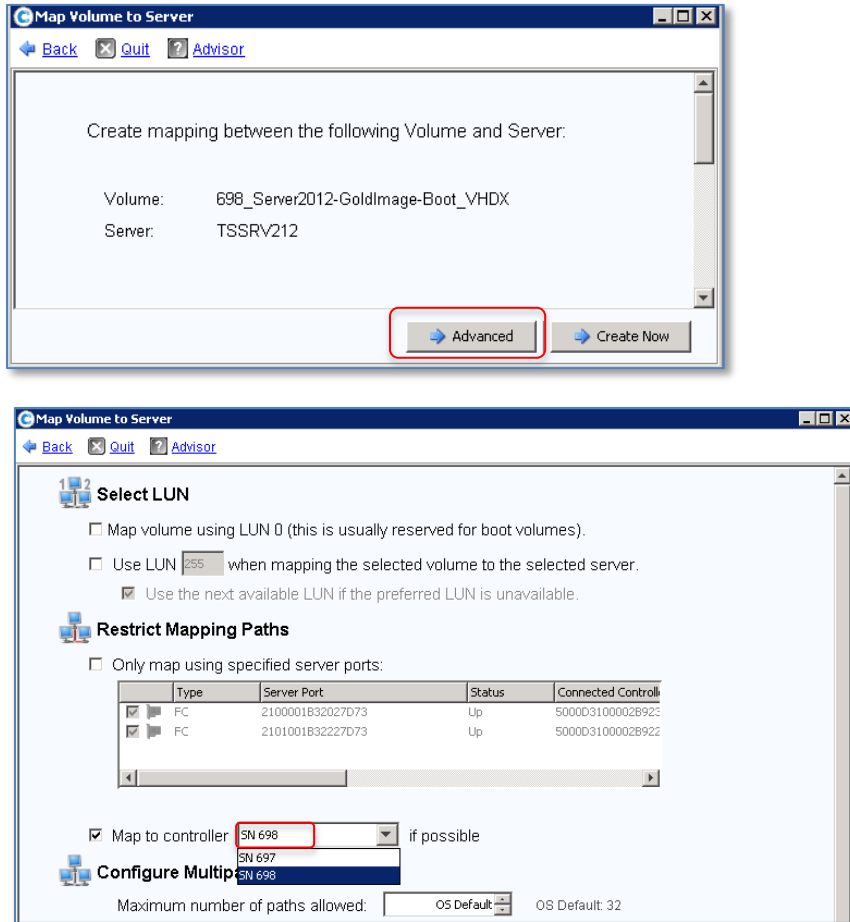


Figure 87. Specify a Controller when Mapping a Gold Image Volume to a Server

- 4) Proceed to map the 2nd gold image volume to a Hyper-V host server. As shown in Figure 87, click on the **Advanced**, button, and then check the box for **Map to controller ___ if possible**.
- 5) From the drop down list, choose the other controller in the pair. In this example, Controller 698 is selected.

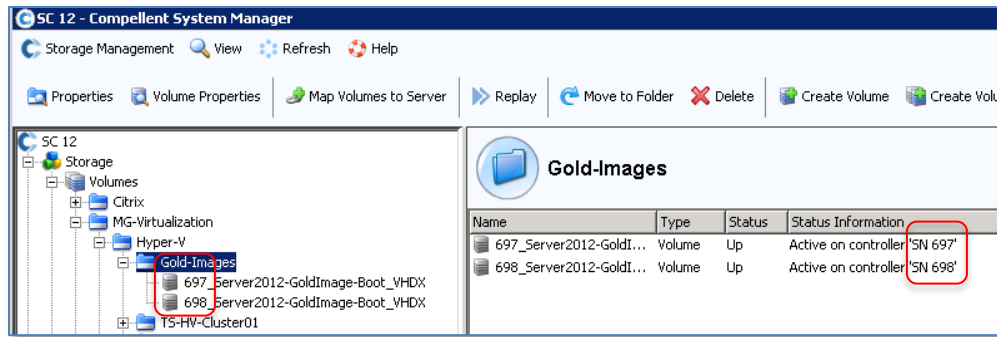


Figure 88. Create a 2nd Gold Image Volume to Balance Load on the Controllers

- 36) Once finished, the two gold image volumes are now mapped to (owned by) different controllers in the pair.
- 37) When deploying new guest VMs from your gold images, alternate between one gold image or the other as needed to help keep the controllers balanced. Use the Tree View under the Replay tab to determine how many guest VMs have been deployed from a gold image.

5.4 Using Gold Images to Deploy Guests to a Hyper-V Cluster

Windows servers assign each volume a unique disk ID (or disk signature). For example, for MBR disks, the disk ID is an 8-character hexadecimal number such as 03E274D2. Each volume mapped to a server must have a unique disk signature.

With standalone Windows servers, the issue of duplicate disk signatures is avoided because standalone servers are able to automatically detect the presence of a duplicate disk signature and automatically change it. This is done without any user intervention.

However, with servers that are clustered, Windows is unable to automatically resolve disk ID conflicts with cluster disks. This issue exists with both Server 2008 R2 Hyper-V and Server 2012 Hyper-V.

Clustering adds a little complexity to deploying Hyper-V guest VMs from gold image volumes on the Dell Compellent Storage Center because of having to deal with duplicate disk signatures.

The first View Volume of a gold image can be presented to a cluster and added as a cluster disk just fine. But the 2nd View Volume, because it has the same disk signature as the first View Volume, creates a disk signature conflict. To resolve the issue, use **Diskpart.exe** to change the disk signature for each View Volume being mapped to a Hyper-V cluster.

For the steps to change the disk signature on a View Volume, please refer to the steps in Section 5.1.4.

The process of creating a gold image to use when deploying Hyper-V guest VMs to a cluster is nearly identical to the steps in Section 5.3 above.

A few additional considerations to keep in mind when using gold images in a clustered environment are as follows:

- 1) The Hyper-V guest VM used to build the gold image can be created on any Hyper-V host server (it does not have to be built on a node in the cluster).
- 2) When a View Volume of the new gold image is mapped to the Cluster, Diskpart must be used to change the disk ID. If the disk ID of the first View Volume is not changed, then when the second View Volume is mapped to the cluster, the disk IDs will conflict.
- 3) Launch **Failover Cluster Manager** and add the new disk to the cluster by right clicking on **Storage** then selecting **Add disk**.
- 4) Create a new Hyper-V guest VM by using **Failover Cluster Manager**.
- 5) When creating the Hyper-V guest VM, select **Store the virtual machine in a different location** and select a path to the newly clustered View Volume.
- 6) When connecting to the virtual hard disk, select **Use an Existing virtual hard disk** and browse to the virtual hard disk at the desired location on the newly mapped volume.
- 7) Complete the configuration of the new Hyper-V guest VM and edit any settings for BIOS, CPU, RAM, networking, etc.
- 8) Power up the new Hyper-V guest VM and verify functionality.

6 Dynamic Capacity

Dell Compellent Thin Provisioning technology (Dynamic Capacity) delivers the highest storage utilization possible. Dynamic Capacity completely separates storage *allocation* (the total size of a volume) from *utilization* (the actual amount of space consumed on the volume by data). This enables administrators to specify larger volume sizes without wasting valuable SAN space with empty volume space that may never get written to. With a Dell Compellent SAN, only the actual data on the volume consumes SAN space. So, a 1 TB volume that has only 10 GB of files will consume only 10 GB of SAN space (plus some overhead for Replays if configured).

Caution: With thin provisioning comes the ability to over-provision the SAN. As long as administrators monitor disk space usage and growth trends carefully, the risk of running a SAN tier completely out of space is minimized.

6.1 Dynamic Capacity with Hyper-V

When a new volume is created on the Dell Compellent Storage Center and presented to a Hyper-V host or guest VM, the volume does not consume any SAN space initially.

After the volume is formatted in Windows and a drive letter or mount point is assigned, a small amount of SAN space is consumed due to the data written to the disk partition table. Only as data pages are actually written to the volume by the host or guest VM will SAN space start to be consumed.

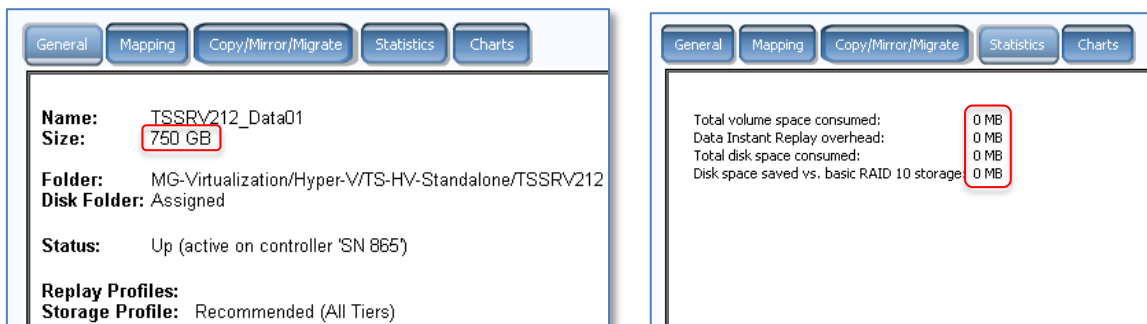


Figure 89. Volume Statistics for a New Volume Showing No Disk Space Used

- 1) In the example shown in Figure 89, a new 750 GB volume mapped to the server TSSRV212 consumes no space initially on the SAN.

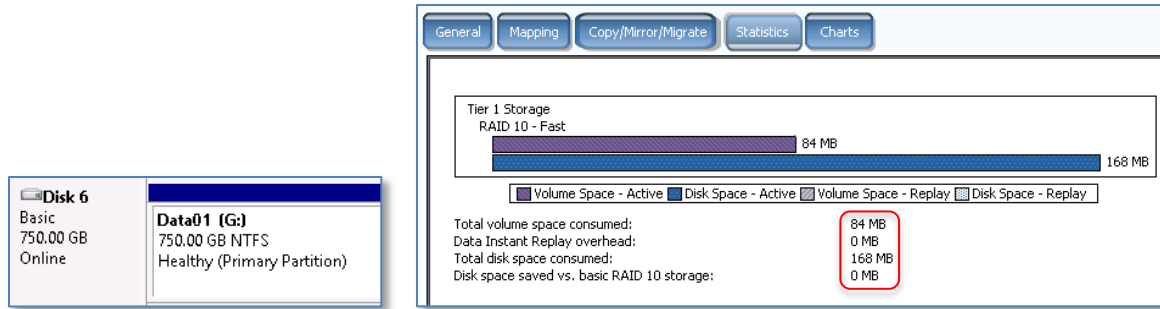


Figure 90. Volume Statistics after a New Volume is Formatted

- 2) After the disk is formatted with windows and assigned a drive letter or mount point, a very small amount of space is used on the SAN. In this example, the 750 GB volume was formatted as an NTFS volume and assigned a drive letter of G.

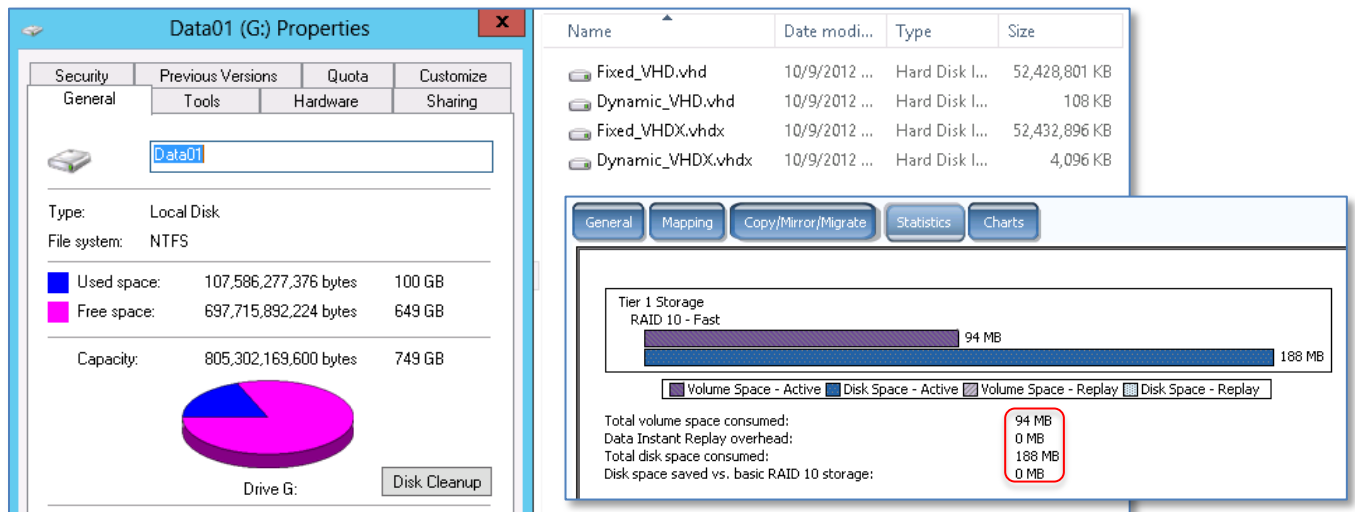


Figure 91. Virtual Hard Disks are Thinly Provisioned

- 3) Virtual hard disk files are also thinly provisioned with Dell Compellent. In this example, 4 virtual hard disks are created on this volume: a 50 GB fixed VHD, a 50 GB dynamic VHD, a 50 GB fixed VHDX, and a 50 GB dynamic VHDX. All four types of virtual hard disk are thinly provisioned on Dell Compellent. So, from the perspective of the windows hosts, about 100 GB of space is now consumed, but the Storage Center reports only 94 MB used, only slightly more space than the 84 MB used after the disk was first formatted and assigned a drive letter.
- 4) Because of Dell Compellent Thin Provisioning, factors other than SAN space utilization (such as performance or copy time) will often determine whether fixed or dynamic VHDs/VHDXs are used for provisioning guest VMs.

7 Data Progression with Hyper-V

The foundation of Dell Compellent's Fluid Data architecture is Automated Tiered Storage, which is the ability to automatically and intelligently move data at the block level to the optimal storage tier and RAID level based on usage and performance, unlike traditional SAN systems.

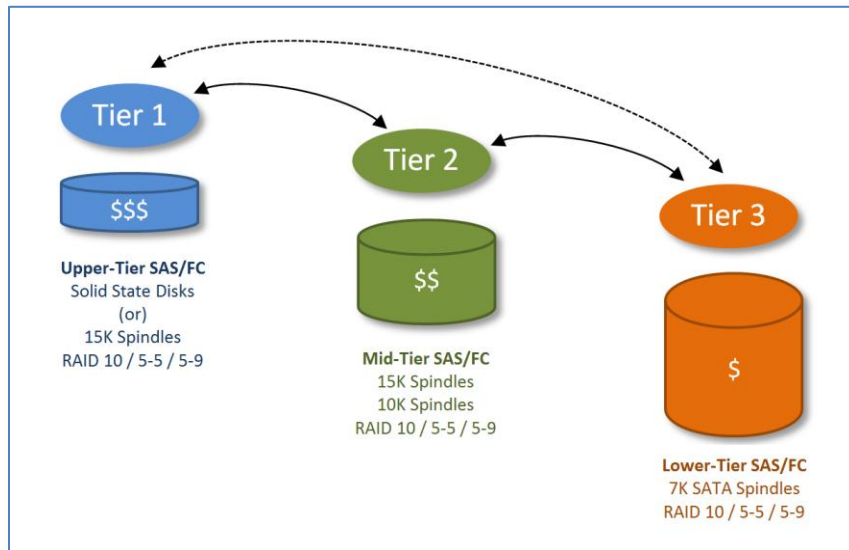


Figure 92. Data Progression

As shown in Figure 92, the highest performing tier also represents the most expensive type of disk, and the lowest performing tier represents the most cost effective disk. In most environments roughly 80% of all data ends up being inactive and therefore, Data Progression can move it automatically to the most cost effective tier (Tier 3).

Hyper-V guest VMs may host a wide variety of operating systems, applications, and associated data. Some data may be inactive or archival in nature. Other data may change frequently and require higher I/O such as transaction logs and page files.

The Dell Compellent system is optimized to automatically move this data at the block level up or down without any administrator intervention, including data residing on Hyper-V hosts and guest VMs with virtual hard disks. With traditional storage solutions, a storage administrator has to manually move data between tiers and RAID levels. With Dell Compellent, the time required for a storage administrator to manage data between different storage tiers and RAID levels is greatly reduced or even eliminated.

With Data Progression, frequently accessed data stays on the highest performing disks (Tier 1-RAID 10). Infrequently accessed data is automatically moved down to lower performance tiers and RAID levels. Data Progression will eventually move static data downward until it reaches the most cost effective tier and RAID level (Tier 3-RAID 5). This provides cost savings

as Tier 3 storage is typically comprised of large-capacity lower-cost SATA disks.

If a new file is created on a Hyper-V guest VM (whether the disk be a VHD/VHDX file, a pass-through disk, or a direct-attached iSCSI or virtual fiber-channel disk), the data is automatically written to the highest storage tier for maximum performance (Tier 1-RAID 10). The longer that file sits without any reads/writes, the further the blocks of data associated with the file will be moved down to lower tiers until it eventually reaches the most cost effective tier (Tier 3-Raid 5). Conversely, a static file that resides in Tier 3 that begins to experience frequent reads will be automatically moved back up towards the highest performing tier (Tier 1-RAID 10).

In most environments, about 80% of all data is static and ends up residing in Tier 3. The remaining 20% of data that is accessed more frequently is kept at Tier 1 or Tier 2 by Data Progression.

7.1 When to Modify Data Progression Settings with Hyper-V

Administrators are able to modify the Data Progression settings for a volume. While this would normally be discouraged, there are some cases where this might be desirable.

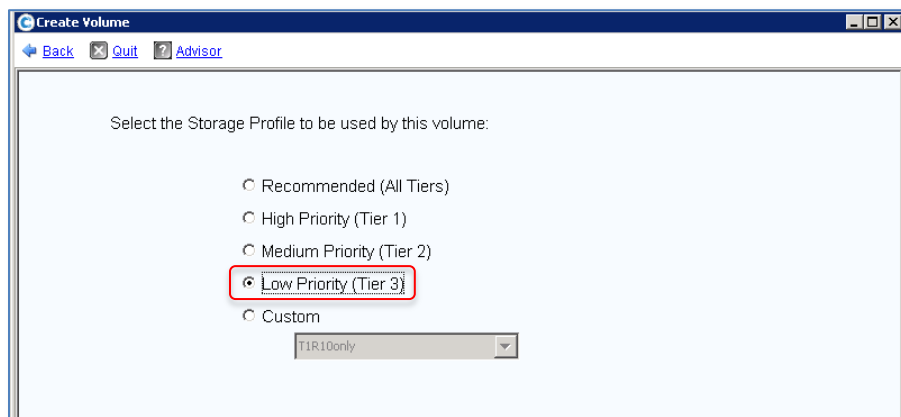


Figure 93. Set a Lower Priority Tier for a New Volume

Example 1: If 1.5 TB of inactive (archival) data needs to be copied from a decommissioned physical server with on-board disk to a new Hyper-V guest VM with a virtual hard disk on a Dell Compellent Storage Center, an administrator may configure the volume hosting the VHD/VHDX file to use Tier 3 only. When the 1.5 TB of data is copied over, it will go straight to Tier 3 instead of consuming Tier 1 SAN space unnecessarily. This also avoids the wait time that would be required for Data Progression to digest this new data and migrate it from Tier 1 down to Tier 3.

Example 2: It might also be necessary to modify Data Progression settings when doing a SAN migration from a legacy SAN to Dell Compellent, or when consolidating a large number of physical servers with on-board disks to Dell Compellent. The danger is that a sudden

influx of a large quantity of migrated data might consume all of the Tier 1-RAID 10 SAN space if there isn't sufficient Tier 1 capacity. If Tier 1 space is consumed, new writes are forced to occur on Tier-2 storage. If Tier 2 storage is consumed, then new writes are forced to occur on Tier 3 storage.

When new writes are forced to occur on a lower tier disk, the SAN may experience noticeable performance degradation. Normally, alert thresholds would notify administrators with enough lead time to allow them to prevent a tier from filling completely up, but in cases where there is a limited amount of Tier 1 capacity, and a large data migration operation is running overnight, Tier 1 capacity might be consumed before an administrator has time to react to a threshold alert.

Since much of the migrated data might not ever need to be on higher tier disks, it may be desirable to push much or all the migrated data to Tier 3 initially. Once the data is migrated, then Data Progression can be configured on those volumes to allow data that needs higher performance to move up to higher performance tiers and RAID levels.

If the inflow of migrated data is slow, then it may be acceptable to allow it to be written at Tier 1-RAID 10 initially (the default) and allow Data Progression migrate the static data down to lower tiers over time.

8 Disk Space Recovery with Hyper-V

8.1 Server 2012 Hyper-V and Trim/Unmap

Support for Trim/Unmap will be introduced with the 6.3.1 release of the Dell Compellent Storage Center OS. With the Trim/Unmap feature, the Dell Compellent Server Agent will no longer be required to recover deleted disk space from server volumes and return it to the SAN's free disk space pool to be used elsewhere. The Dell Compellent Server Agent can still be installed on a 2012 Server, but the disk space recovery feature, since it is no longer needed, will be disabled by default.

Trim/Unmap will be supported with the following types of volumes and disks:

- Dell Compellent SAN volumes mapped to physical 2012 Server hosts using iSCSI or fiber channel, and to guest VMs as pass-through or direct-attached disks using iSCSI or virtual fiber channel:
 - SAN volumes must be "basic" disks and formatted as NTFS (other formats such as FAT and ReFS do not support Trim/Unmap).
- Dell Compellent SAN volumes mapped to Server 2012 Hyper-V nodes as cluster shared volumes (CSVs):
 - CSVs must be "basic" disks and formatted as NTFS (the only supported format for CSVs is NTFS).
- Virtual Hard Disks:
 - The virtual hard disk must be formatted as a .VHDX file (dynamic or fixed). Trim/Unmap is not supported with the .VHD virtual hard disk format.
 - The guest VM OS must support Trim/Unmap. When the guest VM OS is Server 2012, from the perspective of the guest, the VHDX must be a "basic" disk, formatted as NTFS. Trim/Unmap is not supported on a .VHDX when the guest VM's OS is Server 2008 R2 or earlier.

More information about Trim/Unmap support with server 2012 will be included in the *Dell Compellent Best Practices Guide for Windows Server 2012*, due for release in Q1 of 2013.

8.2 Server 2008 R2 Hyper-V and Disk Space Recovery

If 20 GB of data is deleted from an NTFS-formatted volume, Windows "holds on" to the deleted disk space preventing that 20 GB of space from being returned to the free disk space pool on the SAN to be used elsewhere by other servers on the SAN. This is a limitation with how Windows NTFS-formatted volumes work with Server 2008.

Dell Compellent Enterprise Manager, in conjunction with the Dell Compellent Server Agent, provides the ability to recover disk space when files are deleted from NTFS-formatted volumes. This allows the SAN to use this free space elsewhere, thus maximizing thin

provisioning.

There are a couple of ways to leverage disk space recovery with Server 2008 R2 Hyper-V:

- Install the Dell Compellent Server Agent on physical Server 2008 R2 Hyper-V hosts - especially if they have NTFS volumes that experience frequent file deletions.
- Install the Dell Compellent Server Agent on Server 2008 R2 Hyper-V guest VMs with NTFS volumes that are connected either as pass-through disks or direct-attached iSCSI or virtual fiber channel disks.

Disk space recovery will not work on the following types of drives/volumes with Server 2008 R2 Hyper-V:

- Virtual IDE drives presented to guest VMs. Each guest VM allows for up to 4 virtual IDE drives, although typically, only the boot volume is presented to a Hyper-V guest VM as a virtual IDE device.
- VHDs presented to a Hyper-V guest VM. Although VHDs are thinly provisioned, disk space recovery is not possible from a VHD. If free space recovery is highly desirable for a particular guest VM scenario, then present storage (formatted with NTFS) to your guest VMs as pass-through or directly-attached iSCSI or virtual fiber channel disks.
- Even though Microsoft Cluster Shared Volumes (CSVs) are formatted with NTFS, disk space recovery from CSVs is not supported with Server 2008 R2 Hyper-V.

8.3 Configure Disk Space Recovery with Server 2008 R2 Hyper-V

In order to use disk space recovery with Server 2008 R2, the following is required:

- A Dell Compellent Enterprise Manager server must be configured and the Enterprise Manager Data Collector service must be running.
- The Dell Compellent Enterprise Manager Server Agent must be installed on the server. Server 2008 R2 Hyper-V hosts (or Server 2008 R2 guest VMs with direct-attached storage) with SCSI/iSCSI/virtual fiber channel LUNs formatted as NTFS volumes are good candidates for disk space recovery.
- In order for disk space recovery to work, each NTFS volume must have at least one associated Storage Center Replay on the Dell Compellent Storage Center. Configure or assign a Replay schedule to the NTFS volume if no schedule is currently assigned.
- The IP address of each Server 2008 R2 Hyper-V host and guest VM needing disk space recovery must be registered on the Dell Compellent Storage Center.

8.3.1 Disk Space Recovery Example for Server 2008 R2 Hyper-V

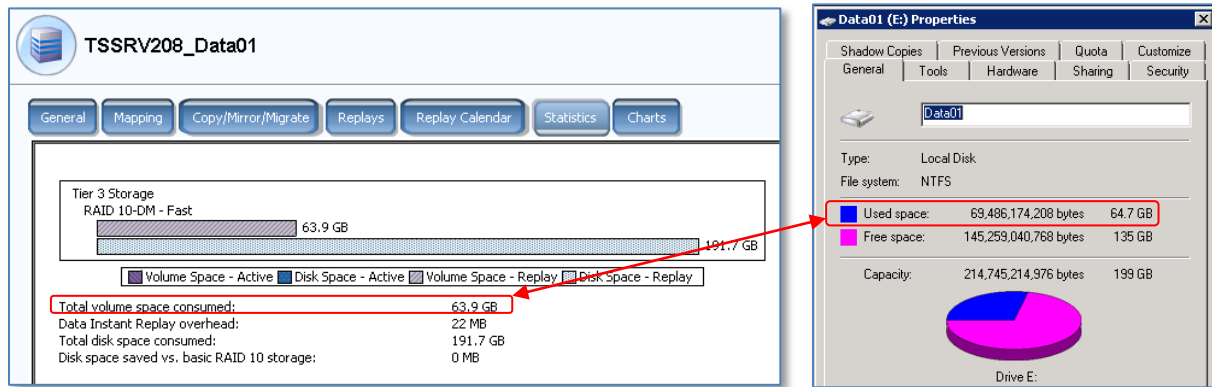


Figure 94. Disk Usage Before Running Disk Space Recovery

- 1) Figure 94 represents a 200 GB volume that is mapped to a Server 2008 R2 Hyper-V guest as a directly-attached iSCSI disk, formatted with NTFS, as the E drive. The volume is currently using about 64 GB of space.

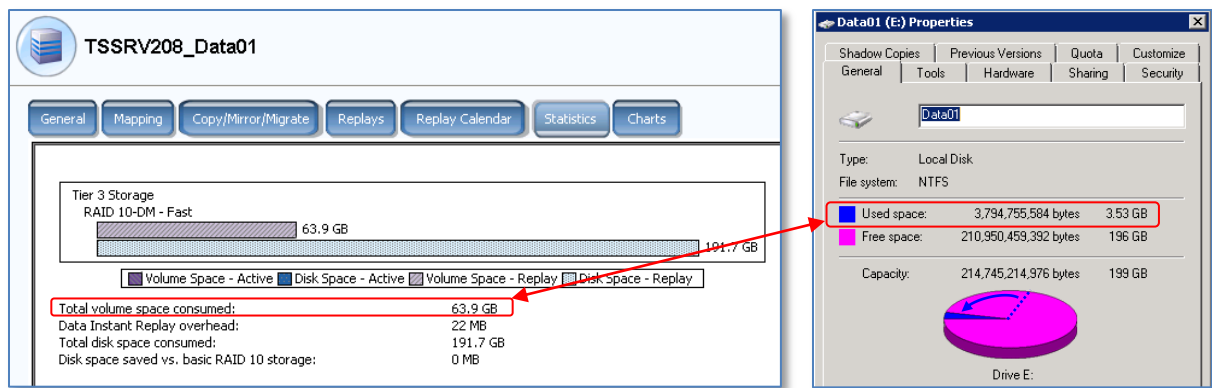


Figure 95. Deleted Space is Not Recovered on the SAN

- 2) Using Windows Explorer on the guest VM, about 60 GB of files were deleted from the E drive. The guest VM shows that the space is now “free” however the SAN still shows that 64 GB of space is consumed. Until disk space recovery runs (this is scheduled to run each evening by default on configured servers), this space will not be available to the SAN to use elsewhere.

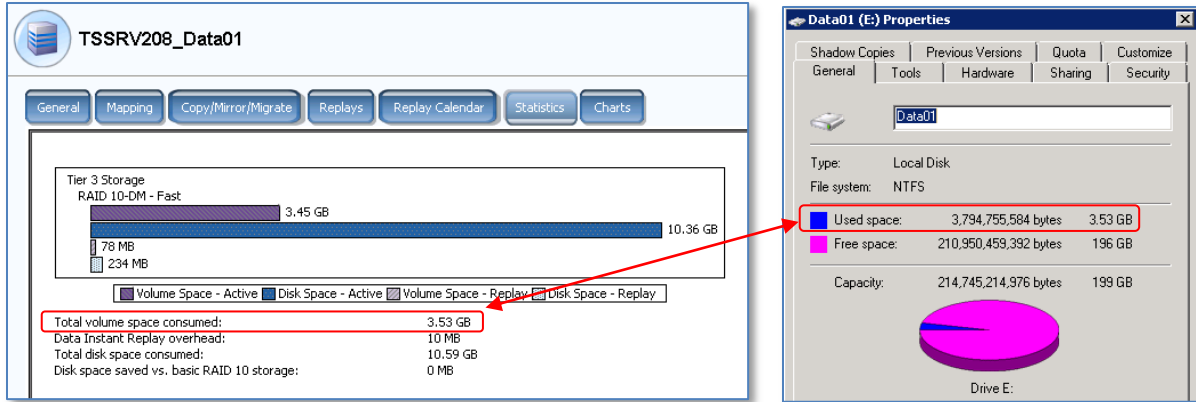


Figure 96. Volume Statistics after Running Disk Space Recovery

- 3) Overnight, disk space recovery runs on this server, and the Dell Compellent Storage Center now shows that the freed-up disk space is now available on the SAN.

9 Hyper-V Performance Tuning

The following sections highlight some performance tuning options for Hyper-V.

9.1 Storage I/O Performance

Hyper-V supports synthetic (virtual) and emulated storage devices in Hyper-V guest VMs but the virtual devices generally offer significantly better throughput, response times and reduced CPU overhead. The exception is if a filter driver can be loaded that reroutes I/O to the emulated storage device. A filter driver is an optional driver (provided by Microsoft or OEM) that modifies the behavior of device hardware to enhance performance.

It is also recommended that virtual SCSI controllers be used instead of virtual IDE controllers for disks that have high disk activity. While the performance difference between virtual IDE and virtual SCSI controllers is generally not a concern for servers with low disk activity, in cases of extremely high disk activity, virtual SCSI is a better choice than virtual IDE as it offers better performance and requires less CPU overhead.

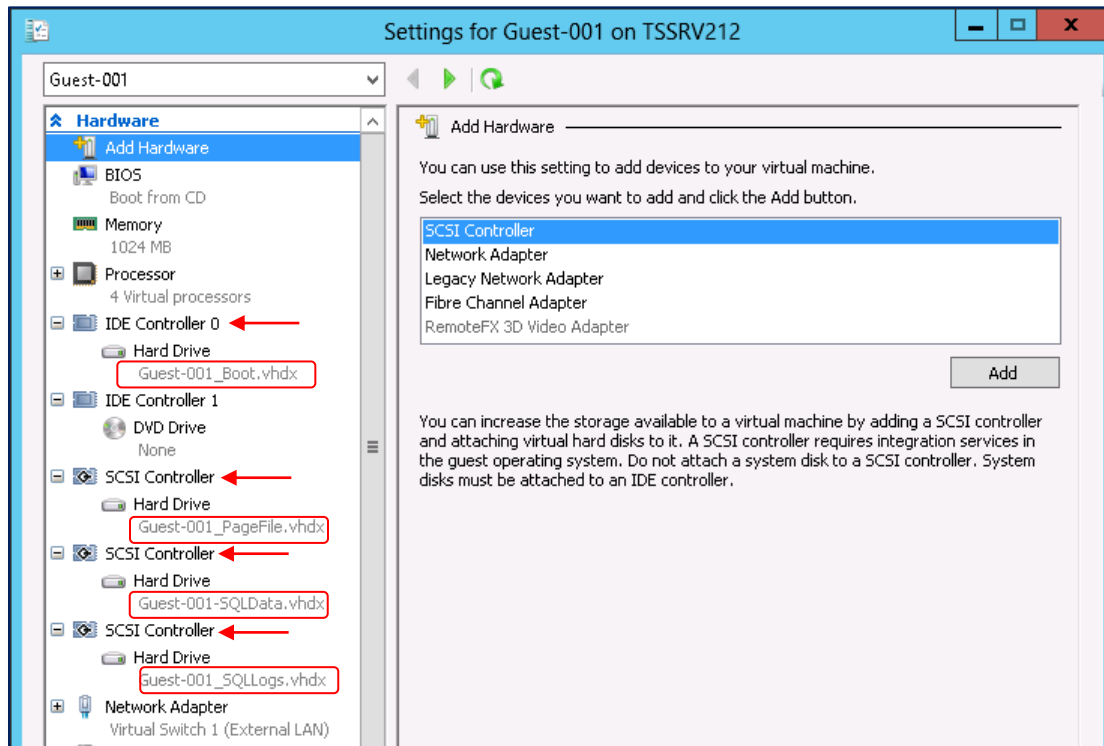


Figure 97. Use Separate Virtual Controllers for High I/O Drives

In addition, if a guest VM has multiple disks requiring high I/O, each virtual hard disk can be associated with its own virtual SCSI controller to further maximize performance. In Figure 97, a Server 2008 R2 guest VM has separate hard drives (VHDXs) for the OS (must be an IDE device), the page file, a SQL database, and for SQL logs.

9.1.1 Block Alignment (Diskpart)

NTFS aligns its metadata and data clusters to a partition boundary by increments of the cluster size (which was selected during file system creation (by default is 4 KB)). In earlier releases of Windows, the partition boundary offset for a specific disk partition could be misaligned as compared to the array disk stripe unit boundaries. This caused requests to be unintentionally split across multiple disks.

To force alignment, use Diskpart.exe at the time that the partition is created. In Windows Server 2008, partitions are automatically offset by 1 MB, which provides good alignment for the power-of-two stripe unit sizes that are typically found in hardware. If the stripe unit size is set to a size that is greater than 1 MB, the alignment issue is much less of a problem because small requests rarely cross large stripe unit boundaries. Note that Windows Server 2008 defaults to a smaller power-of-two offset for small drives.

If alignment is still a problem even with the default offset, use Diskpart.exe to force alternative alignments at the time the partition is created. Since Dell Compellent virtualizes the disk at the block layer, disks that are aligned using Diskpart will show negligible performance differences.

9.1.2 Disable File Last Access Time Check

Windows Server 2003 and earlier Windows operating systems update the “last-accessed” time of a file any time an application opens, reads, or writes to the file. This increases the disk I/O which can negatively impact CPU performance due to the overhead. If applications do not use the “last-accessed” time of files on a server, system administrators should consider setting the following REG_DWORD registry key to a value of 1 to disable these updates.

```
HKLM\System\CurrentControlSet\Control\FileSystem\NTFSDisableLastAccessUpdate
```

By default, newer Windows clients (Vista and newer) and servers (2008 and newer) disable the “last-access” time updates by default. If running Windows 2003 or earlier guest operating systems on Hyper-V, it is recommended that file “last-accessed” be disabled to help optimize disk performance.

9.1.3 I/O Balancer Controls

These are advanced tuning options and should not be used unless there is a specific reason to do so in your environment.

The virtualization stack balances storage I/O streams from different Hyper-V guest VMs so that each guest has similar I/O response times when the system’s I/O bandwidth is saturated. The following registry keys can be used to adjust the balancing algorithm, but note that the virtualization stack already tries to fully use the I/O device’s throughput while providing reasonable balance.

The first path below should be used for storage scenarios, and the second path should be used for networking scenarios:

```
HKLM\System\CurrentControlSet\Services\StorVsp\<Key>
```

```
HKLM\System\CurrentControlSet\Services\VmSwitch\<Key>
```

Both the storage and networking registry paths have the same three REG_DWORD keys:

- IOBalance_Enabled: The balancer is enabled when set to a nonzero value and disabled when set to 0. The default is **enabled** for storage and **disabled** for networking. Enabling the balancing for networking can add significant CPU overhead under some circumstances.
- IOBalance_KeepHwBusyLatencyTarget_Microseconds: This controls how much work (represented by a latency value) the balancer allows to be issued to the hardware before throttling to provide better balance. The default is 83 ms for storage and 2 ms for networking. Lowering this value can improve balance but will reduce some throughput. Lowering it too much may significantly affect overall throughput. Storage systems with high throughput and high latencies can experience increased overall throughput when a higher value is set for this parameter.
- IOBalance_AllowedPercentOverheadDueToFlowSwitching: This controls how much work the balancer issues from a Hyper-V guest before switching to another guest. This setting is primarily for storage where finely interleaving the I/O from different guests can increase the number of disk seeks. The default is 8 percent for both storage and networking.

9.2 Hyper-V Guest Performance

The following sections detail steps that can improve Hyper-V guest performance.

9.2.1 Enlightened I/O Guests

The operating system kernel in Windows clients (Vista SP1 and newer) and Windows Server (2008 and newer) includes feature “enlightenments” that optimize OS performance when the OS is installed as a Hyper-V guest. An enlightened guest means that the OS is virtualization-aware to some degree. The enlightenments work by decreasing the amount of CPU overhead required when an OS is running as a Hyper-V guest VM. In addition, Integration Services (see section 9.2.2) can provide additional enlightenments to help reduce guest VM I/O demands. For optimal guest VM OS performance, Windows Server 2008 or newer is recommended.

9.2.2 Integration Services

The VM integration services include enlightened drivers for the synthetic I/O devices, which significantly reduces CPU overhead for I/O as compared to emulated devices. The latest

version of integration services should be installed on every supported guest VM. These services decrease the CPU demand on the guests and can improve the I/O throughput. Installing the most current version of integration services for the guest OS is normally always the first step in tuning a Hyper-V guest VM for performance.

Because integration services improve the interaction between the guest VM and the Hyper-V host server to ensure optimal performance, it is important to verify that the necessary integration services are installed on the guest VM and that they are the latest version. Without integration services (or if they are outdated) the guest VMs may not run correctly resulting in poor performance and instability.

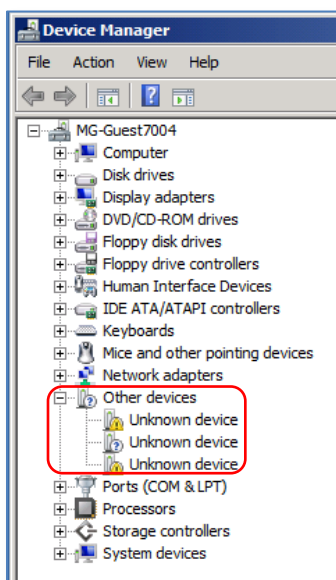


Figure 98. Unknown Guest VM Devices Indicates Missing or Outdated Integration Services

One of the indications that integration services are not installed or outdated on a Windows guest is the presence of unknown devices in Device Manager.

Even though integration services are included natively with Microsoft operating systems such as Windows Server 2008, the integration services may need to be updated depending on the version of Hyper-V host they are running on. For example:

- For a Server 2008 R2 guest VM running on a Server 2008 R2 Hyper-V host, no action should be required to update integration services.
- For a Server 2008 R2 guest VM running on a Server 2012 Hyper-V host, the integration services must be updated.
- If in doubt, verify that the integration services are current by following the steps in section 9.2.3 below.
- For non-Windows guest OS versions that are supported on Hyper-V (e.g. Linux), the latest integration services are available for download from Microsoft.

9.2.3 Installing or Updating Integration Services

In the following example, the integration services will be updated on a new install of Server 2008 R2 running as a guest VM on a Server 2012 Hyper-V host.

- 1) Connect to the Server 2008 R2 guest VM using Hyper-V Manager or Failover Cluster Manager on the host server.

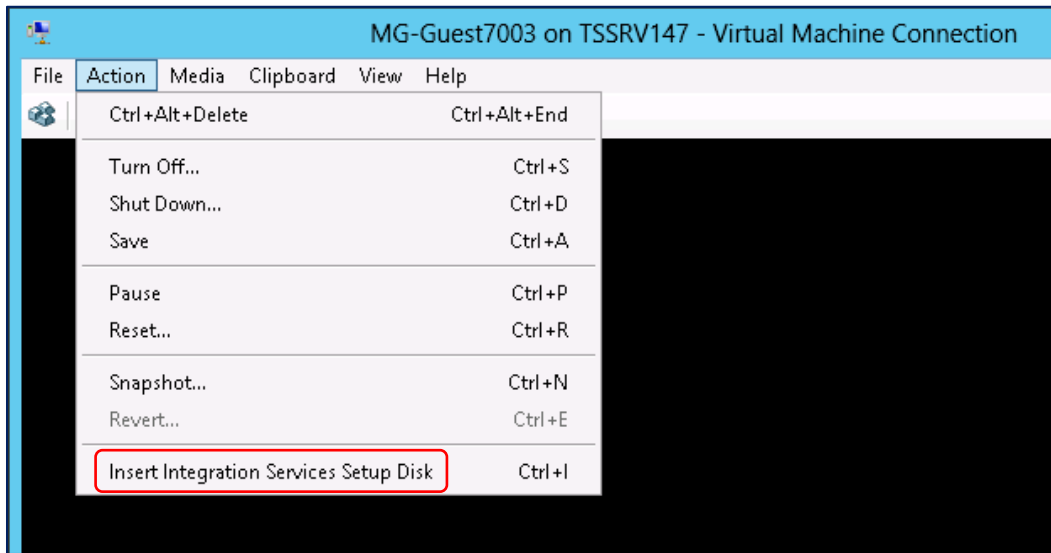


Figure 99. Install Integration Services

- 2) After connecting to the server, from the **Action** menu drop-down, select **Insert Integration Services Setup Disk** as shown in Figure 99.

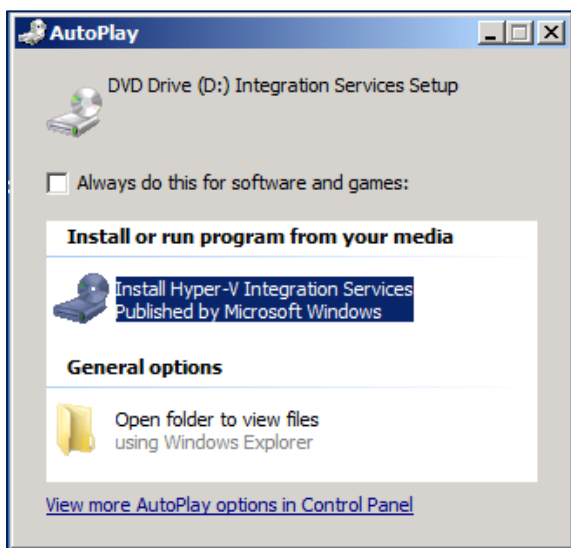


Figure 100. AutoPlay Detects the Integration Services Installer

- 3) After a few seconds, AutoPlay should present the integration services installer. Click

on it to launch the installer. If no AutoPlay option appears, browse to the mounted DVD drive on the guest VM and double click on it to launch the installer.

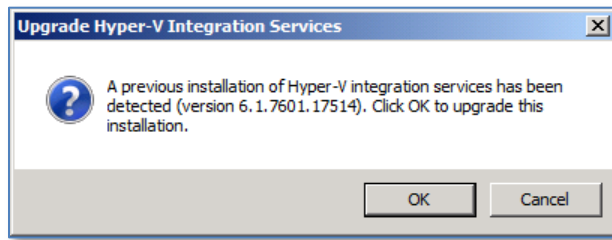


Figure 101. Upgrade Hyper-V Integration Services Dialog Box

- 4) Review the message and click on **OK** to continue the install.
- 5) Once the install has finished, reboot the guest VM.
- 6) After rebooting, verify that Device Manager on the guest VM is clean. It should not contain any unknown devices.

9.2.4 Background Activity

Minimizing the background activity in idle Hyper-V guest VMs frees up CPU cycles that can be used elsewhere by other guests, or to reduce power consumption. Windows Hyper-V guests will typically use less than 1 percent of a CPU core when they are idle.

The following are several best practices recommendations for minimizing the background noise on a Windows Server or Client that is running as a Hyper-V guest VM to minimize CPU usage:

- Install the latest version of the integration services for the guest VM's OS.
- Disable the screen saver or select a blank screen saver.
- Remove or disable any unused hardware devices such as the COM port or CD-ROM/DVD (or at least disconnect any mounted ISO media).
- Unless required by an application or process, log off the guest VM. An inactive logon will consume a small amount of resources unnecessarily.
- Use newer versions of Windows Server/Client (Server 2008 or newer; Windows 7 or newer) as these operating systems are more "enlightened" (virtualization aware) allowing for more efficient CPU usage.
- Disable, throttle, or stagger scheduled tasks, periodic maintenance, or similar recurring tasks or activities such as backups, virus scans, automatic updates, patching, and defragmentation as appropriate.
- Disable any Windows or application services enabled by default that are not needed.
- Tune server/client applications to reduce unnecessary periodic activity.
- Disable any unnecessary background services such as Windows Search.
- Use Control Panel to disable the Aero Glass desktop experience with Vista and Windows 7 clients.

9.2.5 Windows Page File

Windows Server and Client OS installations place the page file on the boot volume by default, and automatically manage it without user intervention. While Microsoft typically recommends moving the page file to a separate dedicated volume for a Server OS to maximize performance, there are some practical considerations that come into play as to when and if this is really necessary. Ultimately, the decision will be a function of factors and variables that are unique to each environment.

If a Hyper-V guest VM is running a Windows Server or Client OS with low memory and disk I/O demands, and can perform acceptably well with the default configuration, adding a 2nd drive and moving the page file (which adds set-up time and complexity) may not make any sense. The general best practices recommendation is to keep the guest VM virtualization design as simple as possible, and vary from the defaults only when necessary for design, performance, or recoverability reasons.

For Hyper-V guest VMs running Windows Servers or Clients that host memory or disk intensive applications (or if required by a particular application's best practices design guide), separating the page file to its own volume may be necessary to insure optimal performance.

There are also some definite advantages to placing a page file on a separate volume from the Dell Compellent perspective. The below reasons in and of themselves may not be sufficiently advantageous to make it worth the while to vary from the default page file design, but should be considered as part of the overall page file strategy.

- Moving the page file to a separate dedicated volume reduces the amount of data that is changing on the system (boot) volume. This helps reduce the size of Dell Compellent Replays for the boot volume, consuming less SAN space.
- Volumes or virtual hard disks dedicated to page files typically do not require Replays or snapshots taken for recovery purposes, and therefore don't need to be replicated to a remote Storage Center. This is especially beneficial in cases where there is limited bandwidth for replication of volumes and Replays between Storage Centers.
- Volumes containing only page files can also be excluded from Data Progression. Since page file data is constantly changing, it benefits from staying at Tier 1-RAID 10 for maximum performance.

10 Boot from SAN for Hyper-V

In some cases, such as with blade servers that do not have internal disk drives, booting from SAN is the only option. Many physical Hyper-V hosts have internal drives providing for the ability to boot locally or from the SAN. Deciding which option to use is dependent on a number of considerations that are unique to each environment. There are some advantages and disadvantages to each option:

10.1 Advantages of Booting from Dell Compellent SAN

- Ability to take Dell Compellent Snapshots (Data Instant Replays) of boot volumes which allows for quick recovery.
- Ability to replicate boot volumes to a remote location for disaster recovery purposes.
- Ability to create and leverage gold images of a Windows 2008 or 2012 Hyper-V host server boot volume to quickly provision new Hyper-V host servers.
- If Dell Compellent Replays (with Remote Replication) are the strategy for short and long term backup and recovery, when boot volumes are on the SAN, there is no need for other backup software to back up data on those volumes.

10.2 Advantages of Booting from Local Disk

- In most cases, down time is avoided due to configuring redundancy for the SAN (dual controllers), data paths (dual fabrics), network paths (dual switched networks) and dual power domains, each with UPS backup. In the rare event (despite redundancies) that an event prevents host server access to the SAN, Hyper-V host servers with local boot disks will stay on line. However, any data disks that are SAN-based will be affected.
- If booting from local disk there is no need to purchase fiber channel or iSCSI HBA adapters. For data volumes, the software iSCSI initiator can be used with standard LAN-on-motherboard (LOM) or PCI-based NIC cards to map to data volumes on the SAN.
- The Dell Compellent Enterprise Manager and Data Collector servers should be on a server with local boot disk so that in case of service events affecting the SAN, administrators can still access the Enterprise Manager GUI.
- It is also recommended to have a backup Active Directory (AD) domain controller configured to boot from local disk to help ensure domain functionality for AD-dependent resources in cases where the SAN needs to be completely off line for maintenance.

10.2.1 Domain Controllers on Hyper-V Clusters

Do not configure your AD domain controllers as Hyper-V guest VMs on Hyper-V Clusters that depend on those AD servers for Cluster Services authentication.

This is because there is a boot-order dependency that requires at least one AD server to be on line before cluster services can authenticate and start. If those dependent AD servers are down because they are on the affected cluster, the cluster then has no way to authenticate and start.

This unfortunate scenario will result in unexpected down time. To recover, manually recover the guest VM running as an AD domain controller to a standalone hyper-V host outside of the cluster. For example, map View Volumes containing the guest VM's virtual hard disks to a standalone host outside of the affected cluster, and manually recover the guest VM there.

10.3 Configuring Hyper-V Hosts to Boot from SAN

Physical Hyper-V servers can be configured to boot from SAN using physical fiber channel or iSCSI host bus adapters (HBAs). An example of each type of configuration will be shown below.

10.3.1 Configure Boot from SAN with Fiber Channel HBAs

In this example, a rack mount physical Hyper-V host server with a QLE2562 (dual port) 8 GB fiber channel HBA will be mapped to a dual-controller Dell Compellent Storage Center, boot from SAN. Each port of the HBA will be presented to a separate fabric to insure redundancy and to provide for multipath I/O (MPIO) load balancing.

In this example, the Hyper-V host server will be booting from a sysprepped gold image that already has the MPIO feature installed and configured. Therefore both HBA paths can be enabled and configured at the same time. If staging a new Hyper-V host from scratch that has not had the MPIO feature installed, then modify the below steps to enable only one HBA path initially. Once the OS and the MPIO feature are installed, then enable the 2nd HBA path and proceed to configure MPIO.

For more information on enabling and configuring MPIO settings for Server 2008 and Server 2012 Hyper-V, please refer to Section 4 of this document.

- 1) Ensure that the fiber channel switches, server HBA ports and cabling are configured properly to allow both HBAs in the physical Hyper-V host to advertise their world wide names (WWNs) to the fabric so that zoning steps can be completed.

A screenshot of a QLogic BIOS settings screen. The text is white on a black background. It reads: 'QLE2562 PCI3.0 Fibre Channel ROM BIOS Version 3.00', 'Copyright (C) QLogic Corporation 1993-2010. All rights reserved.', 'www.qlogic.com', 'Press <CTRL-Q> or <ALT-Q> for Fast!UTIL', and 'Firmware Version 5.04.01'.

```
QLE2562 PCI3.0 Fibre Channel ROM BIOS Version 3.00
Copyright (C) QLogic Corporation 1993-2010. All rights reserved.
www.qlogic.com

Press <CTRL-Q> or <ALT-Q> for Fast!UTIL
Firmware Version 5.04.01
```

Figure 102. Select CTRL-Q to Access the QLogic BIOS Settings

- 2) Boot the physical server with the QLogic HBA card and during the boot process, press **Ctrl -Q** when prompted to enter the HBA **Fast!UTIL** configuration tool.

Select Host Adapter					
Adapter Type	Address	Slot	Bus	Device	Function
QLE2562	4000	02	13	00	0
QLE2562	4400	02	13	00	1

Figure 103. QLogic Host Adapter Main Screen

- 3) Since this is a dual-port QLogic HBA adapter, it displays two HBAs on the **Select Host Adapter** Screen. In this example, the top HBA with address 4000 will be referred to as **HBA1**, and the bottom HBA with address 4400 will be referred to as **HBA2**. These HBAs will be associated with separate fabrics to provide redundancy.
- 4) With the top HBA highlighted (**HBA1**), press **Enter**.
- 5) With **Configuration Settings** highlighted, press **Enter** again.



Figure 104. Reset the HBA to Default Settings

- 6) Arrow down to **Restore Default Settings** and press **Enter**. While restoring the default settings is not required, it is recommended to clear any previous HBA settings that might interfere with a successful configuration.
- 7) With **Adapter Settings** highlighted, press **Enter** again.

Adapter Settings	
BIOS Address:	70000
BIOS Revision:	3.00
Adapter Serial Number:	LFD1144W76894
Interrupt Level:	7
Adapter Port Name:	21000024FF3E8A0C
Host Adapter BIOS:	Enabled
Frame Size:	2048
Loop Reset Delay:	5
Adapter Hard Loop ID:	Disabled
Hard Loop ID:	0
Spinup Delay:	Disabled
Connection Options:	1
Fibre Channel Tape Support:	Enabled
Data Rate:	2

Figure 105. Set the HBA Adapter Settings

- 8) Modify the **Adapter Settings** as shown in figure 105:
 - **Enable** the **Host Adapter BIOS**
 - Be sure to set **Connection Options** to **1** (Point to Point only)
 - Leave all other options set to the defaults (as shown)

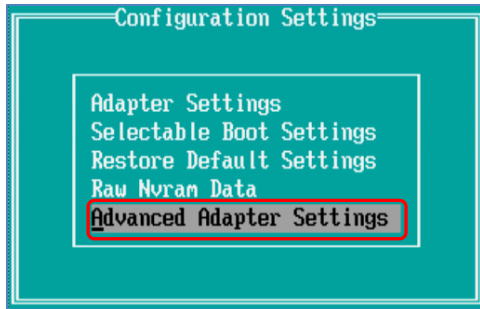


Figure 106. Advanced HBA Settings

- When finished, press the **Escape** key once, and then select **Advanced Adapter Settings** and press **Enter**.




Figure 107. Set Advanced HBA Settings

- Configure the **Advanced Settings** as show in figure 107:
 - For **Execution Throttle** (queue depth), set a higher or lower number than **256** if desired for your server environment
 - Set **Luns per Target** as desired (set to **256** in this example)
 - Set **Enable LIP Reset**, **Enable LIP Full Login**, and **Enable Target Reset** to **Yes**
 - Set **Login Retry Count** to **60** (seconds)
 - Set **Port Down Retry Timeout** to **60** (seconds)
 - Set **Link down Timeout** to **30** (seconds)
 - Leave the other options set to the defaults as shown

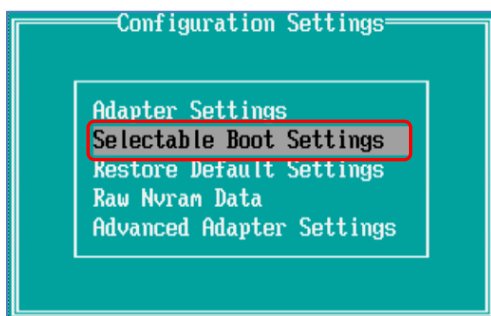


Figure 108. Selectable Boot Settings

11) Press **Escape** one time and then highlight **Selectable Boot Settings** and press **Enter**.




Figure 109. Enable Selectable Boot Settings

12) Set **Selectable Boot** to **Enabled**. Make sure that all boot ports are cleared (so they show all zeros).




Figure 110. Save HBA Configuration Changes

13) Then press **Escape** twice, and when prompted to **Save Changes**, do so.



Figure 111. Select Another HBA

14) Arrow down and choose **Select Host Adapter**, and then press **Enter**.




Figure 112. Select the 2nd HBA

15) Repeat steps 4 through 13 above to configure the 2nd HBA (**HBA2**) and enable it, and save the changes.

16) Reboot the server, and press **Ctrl -Q** when prompted to enter the HBA **Fast!UTIL** configuration tool again.

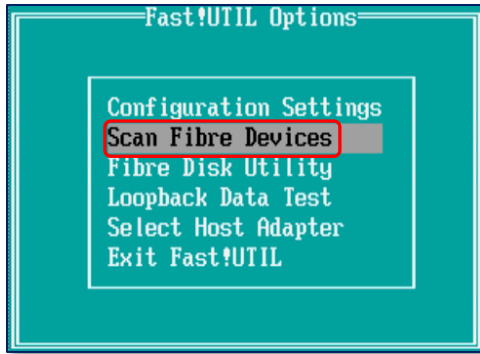


Figure 113. Scan Fiber Devices

17) Highlight **HBA1**, press **Enter**, arrow down to **Scan Fiber Devices** and press **Enter**.

ID	Vendor	Product	Rev	Port Name	Port ID
0		No device present			
1		No device present			
2		No device present			
3		No device present			
4		No device present			
5		No device present			
6		No device present			
7		No device present			

Figure 114. No Devices Present

18) No active ports will be shown (yet) but the HBA card can now advertise its World Wide Name (WWN) to the fabric so that zoning steps can be completed.

19) Repeat step 17 for **HBA2**.

20) At this point, leave the server powered on sitting at the QLogic **Fast!UTIL** screen and complete any fabric configuration steps to allow **HBA1** and **HBA2** to see the Dell Compellent Storage Center. This typically involves creating a zone on each fabric to associate the desired Storage Center with the WWN for the desired server HBA (or associated fiber switch port).

21) Once the fabric changes have been completed, the WWNs for **HBA1** and **HBA2** should now be visible to the Storage Center. Log on to Storage Center Manager and under Servers, create a server folder to hold the new server object.

22) Right click on this folder and select **Create Server**.

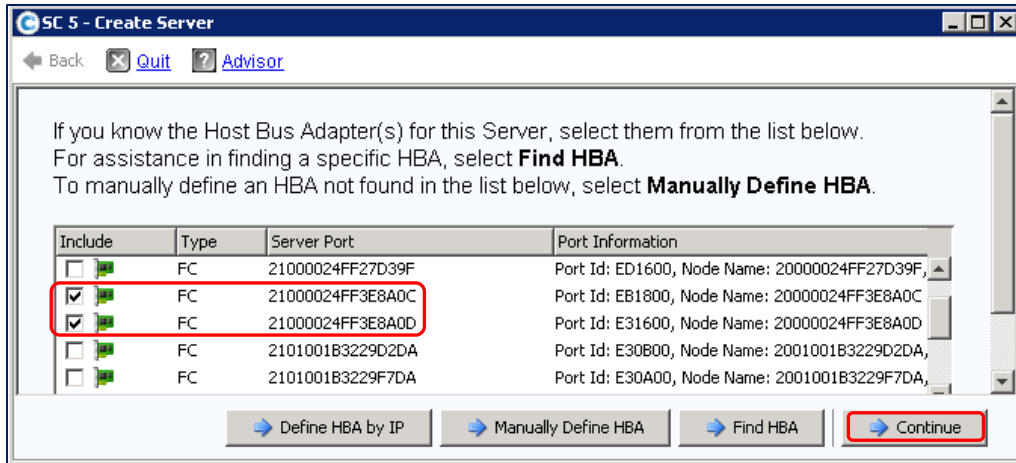


Figure 115. Associate the HBAs to a Server Object on the Dell Compellent Storage Center

23) The WWNs of both **HBA1** and **HBA2** should now show as available in the listing of available HBAs. In the **Include** column, check the box for **HBA1** and **HBA2** and then click on **Continue**.

Note: If the HBAs are not listed as shown in Figure 115, then double-check your fabric configuration, HBA settings and cabling. Repeat **Rescan Fibre Devices** using the **Fast!UTIL** utility until the HBAs are visible in the list. While it is possible to manually define HBAs if they are not listed, the best practices way is to choose them from the list as it removes all doubt that the server’s HBAs can “see” the Storage Center successfully.

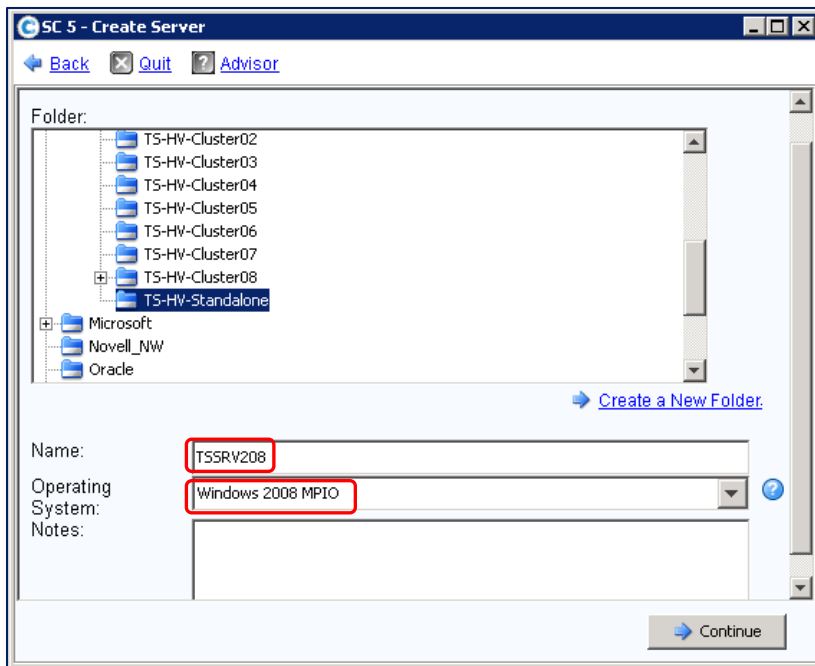


Figure 116. Configure Server Settings

24) Provide a name for the server, and from the **Operating System** drop down list, choose the appropriate operating system (**Windows 2008 MPIO** in this example) and

then click on **Continue**, and then **Create Now**.

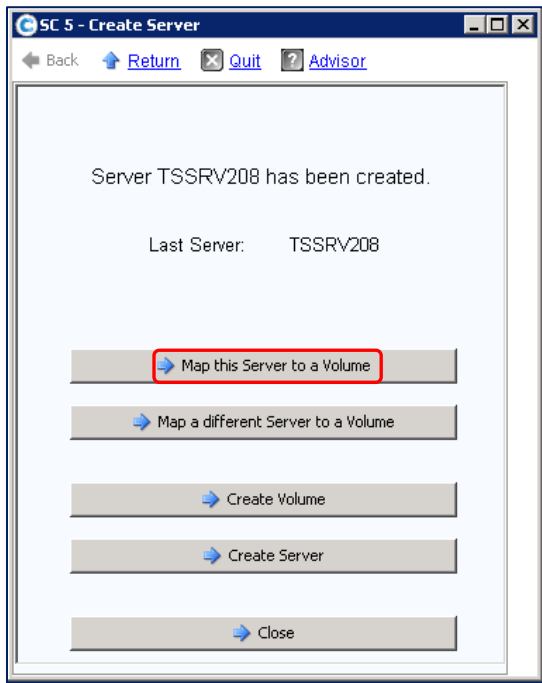


Figure 117. Map the Server to a Volume

25) Once the server has been created, it needs to be mapped to a boot volume. Do one of the following:

- Select **Map this Server to a Volume** if a boot volume already exists (as in this example)
- Select **Create Volume** if a boot volume needs to be created

26) In this example, since the boot volume already exists, **Map this Server to a Volume** is selected as shown in Figure 117.

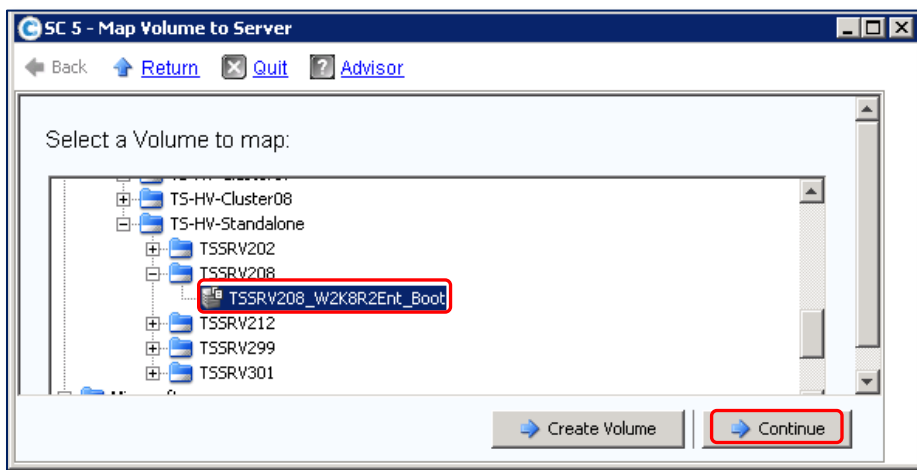


Figure 118. Map the Server to a Boot Volume

27) Select the desired boot volume, and then click on **Continue**.

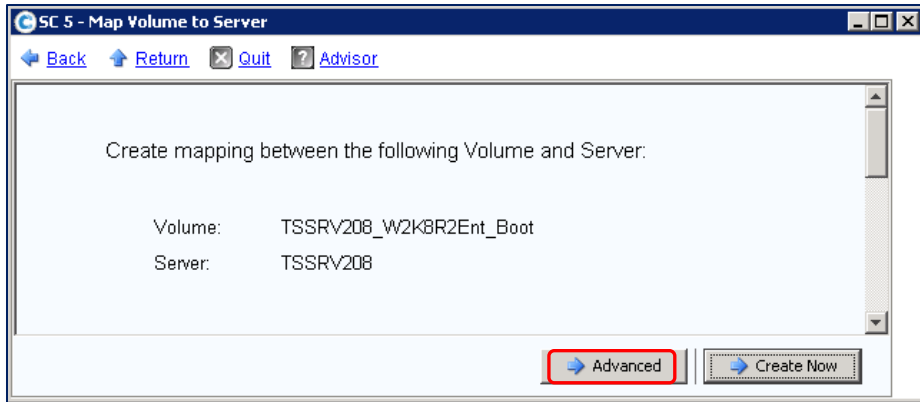


Figure 119. Select the Advanced Button to Configure LUN Settings

28) Click on the **Advanced** button.

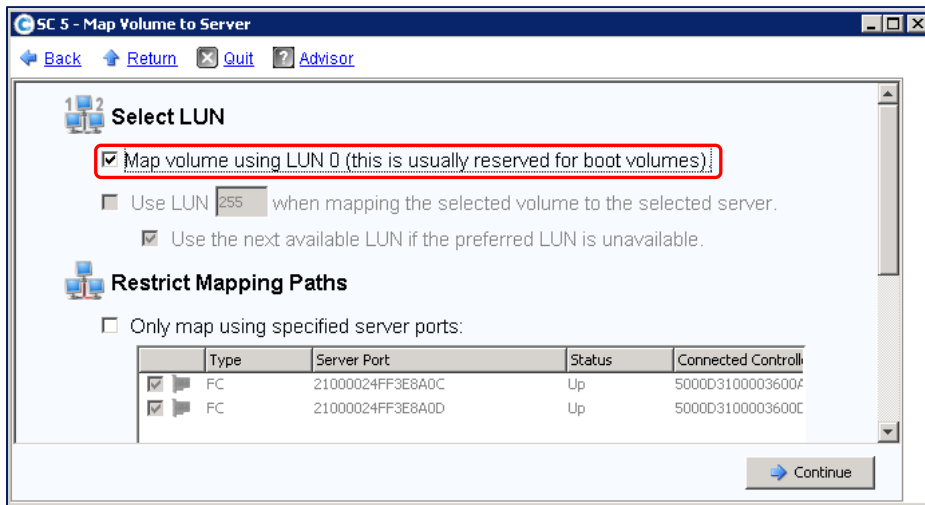


Figure 120. Set the Boot Volume to use LUN 0

29) Because this is a boot volume, check the box for **Map volume using LUN 0**. Then click on **Continue**, **Create Now**, then on **Close**.

30) Return to the QLogic **Fast!UTIL** configuration screen on the physical host server.

31) Select **HBA1** and press **Enter**.

32) Select **Scan Fibre Devices** and press Enter.




Figure 121. A Scan Now Shows an Active Boot Path for HBA1

33) The scan results should now show one or more active ports as shown in Figure 121.

34) Repeat with **HBA2** to ensure it can also see one or more active ports.

35) Select **HBA1** again and press **Enter**. Under **Configuration Settings**, choose **Selectable Boot Settings** and press **Enter**.

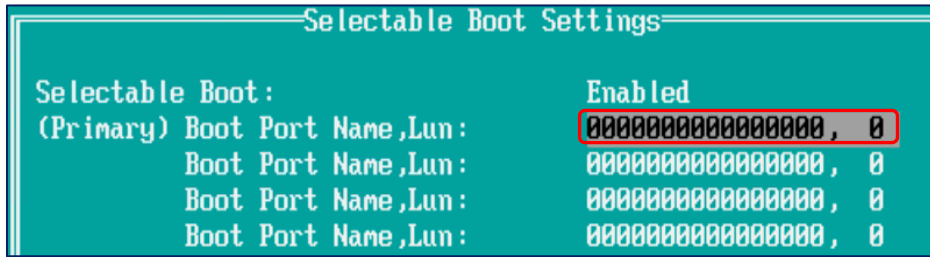


Figure 122. Highlight the 1st Boot Port

36) Arrow down to the first **Boot Port** (which should show all zeros) and press **Enter**.




Figure 123. Select the Active Path

37) After a brief pause, at least one active port should be listed as shown in Figure 123. In this example, there is only one active port for each HBA. Arrow down to the first active port and press **Enter**.

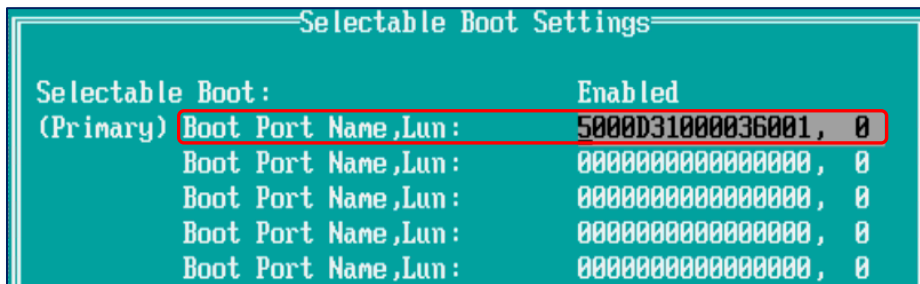


Figure 124. Active Boot Port Now Listed for HBA1

38) The active boot port should now be populated with a WWN as shown in Figure 124.

39) If there are more active ports (as listed in Figure 123), arrow down to the next **Boot Port**, press **Enter**, and arrow down and highlight the next active port listed, and press **Enter**.

40) When finished adding active boot ports for this HBA, press **Escape** twice and save the changes. In this example, there is only one active boot port listed for **HBA1**.

41) Repeat the same process for **HBA2**, following steps 35 - 40, until all its active boot ports are listed under **Selectable Boot Settings**.

42) Save the HBA changes and reboot the server.

```

Press <CTRL-Q> or <ALT-Q> for Fast!UTIL
ISP24xx Firmware Version 4.00.30

Device Device Adapter Port Lun Vendor Product Product
Number Type Number ID Number ID ID Revision
Disk 0 C70800 0 COMPELNT Compellent Vol 0601
ROM BIOS Installed

```

Figure 125. Dell Compellent Boot Disk Listed at Server Startup

43) As the server starts to boot, a Dell Compellent Disk (with a LUN of 0) should be listed as shown in Figure 125 before the OS starts to load.

Note: If the physical host server has an on-board RAID controller, it may be necessary to disable the onboard controller in the host server’s BIOS, and/or configure the boot order priority so that the QLogic HBA is listed before the onboard RAID controller. Otherwise the server may ignore the QLogic HBA boot path and attempt to boot from local disk.

44) Since the boot volume in this example is a sysprepped server gold image that already has the MPIO feature installed with MPIO configured, no other action is required.

Once the OS is loaded, verify (and adjust if desired) the MPIO settings.

45) If an OS is being installed from scratch, modify the above steps as follows:

- a) Install the OS with only a single HBA path and single port enabled.
- b) Once the OS is installed, install the Multipath MPIO feature.
- c) Reboot the server and use the QLogic Fast!UTIL utility to enable the 2nd HBA path.
- d) Complete any fabric changes for the 2nd path.
- e) Add the 2nd HBA to the server in Dell Compellent.
- f) Boot the server and complete MPIO configuration.
- g) For more information on enabling and configuring MPIO settings, please refer to Section 4 of this document.

10.3.2 Configure a QLogic iSCSI HBA Card to Boot from SAN

The below steps show how to configure a Hyper-V host server to boot from SAN using a QLogic iSCSI HBA card. In the below example, a single path will be configured. If MPIO is desired, then once the OS has been loaded, enable the MPIO feature (see Section 4) and repeat the below steps to enable a 2nd iSCSI path, and then configure MPIO settings.

- 1) Install the QLogic iSCSI HBA into the physical server, complete any required switch configuration, connect any required cabling, then power on the server.

```
QLogic Corporation
QLE4062C iSCSI ROM BIOS Version 1.13
Copyright (C) QLogic Corporation 1993-2007. All rights reserved.
www.qlogic.com
```

Press <CTRL-Q> for Fast!UTIL

```
<CTRL-Q> Detected, Initialization in progress, Please wait...
```

Figure 126. Press CTRL-Q to Access the QLogic iSCSI Fast!UTIL utility

- 2) Press **Ctrl-Q** when prompted to enter the QLogic **Fast!UTIL** utility.
- 3) If there are multiple iSCSI adapters listed, select the desired adapter to configure.

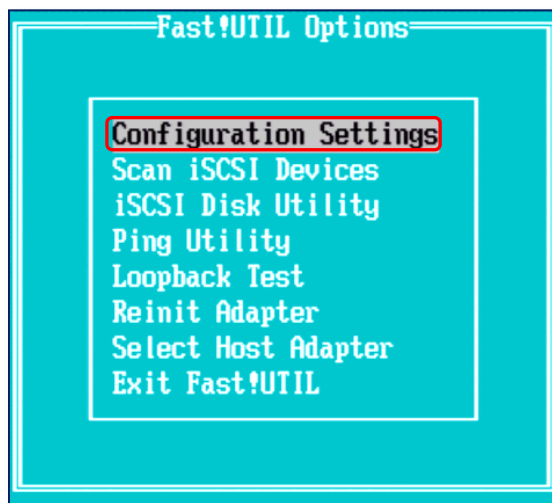


Figure 127. Select QLogic iSCSI Configuration Settings

- 4) From the **Fast!UTIL Options** menu select **Configuration Settings**.

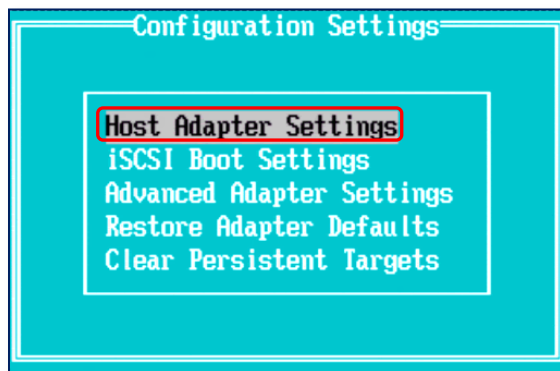


Figure 128. Select QLogic Host Adapter Settings

- 5) From the **Configuration Settings** window, select **Host Adapter Settings**.



Figure 129. Select QLogic Initiator IP Settings

- 6) On the **Host Adapter Settings** menu, select **Initiator IP Settings**.

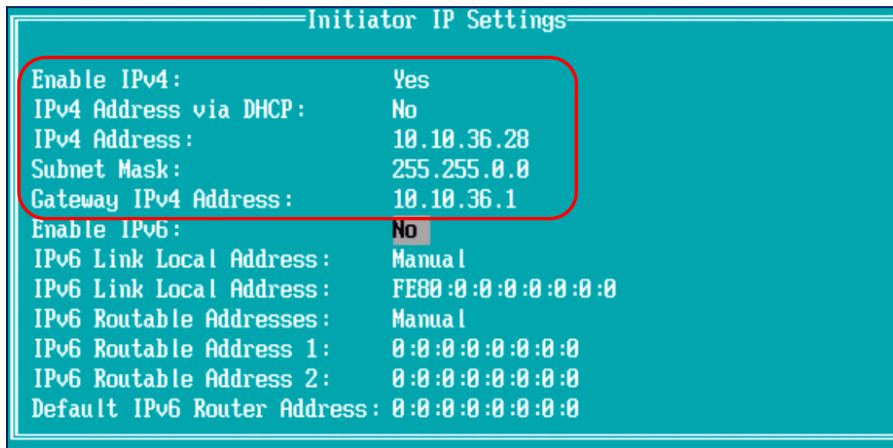


Figure 130. QLogic Initiator IP Settings

- 7) On the **Initiator IP Settings** menu, enter the appropriate information for your configuration such as **IP Address**, **Mask**, and **Gateway**. In this example, IPv4 is used.
- 8) When finished, press **Escape** to go back to the **Host Adapter Settings** menu.

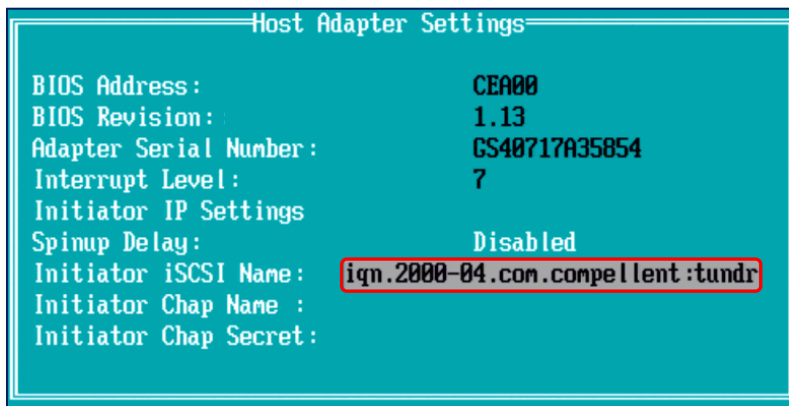


Figure 131. QLogic Initiator iSCSI Name

- 9) Select **Initiator iSCSI Name**.

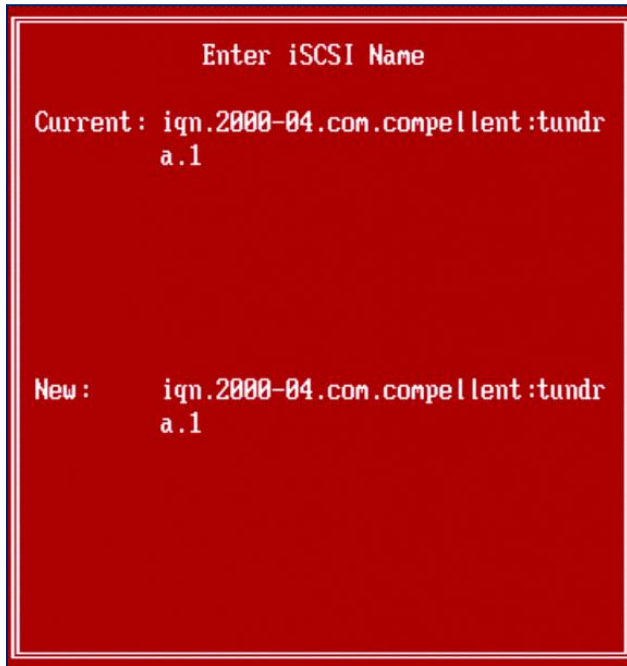


Figure 132. Enter iSCSI Name

- 10) Enter an appropriate **iSCSI name** for the connection.
- 11) Hit **Escape** to go back to the **Host Adapter Settings** menu.
- 12) Make sure **Spinup delay is Disabled**.
- 13) Hit **Esc** to return to the **Configuration Settings** menu.
- 14) Select **iSCSI Boot Settings**.

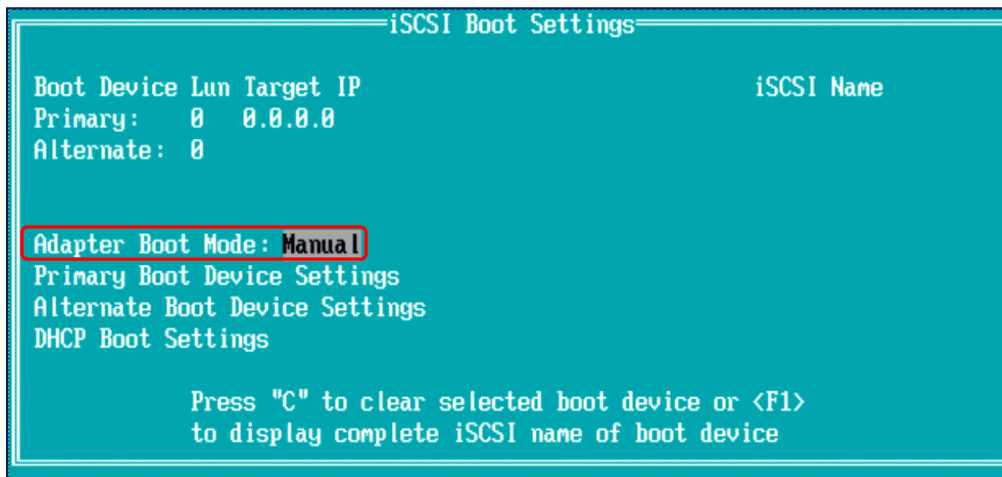


Figure 133. iSCSI Boot Settings

- 15) On the **iSCSI Boot Settings** screen select **Adapter Boot Mode** and set it to **Manual**.

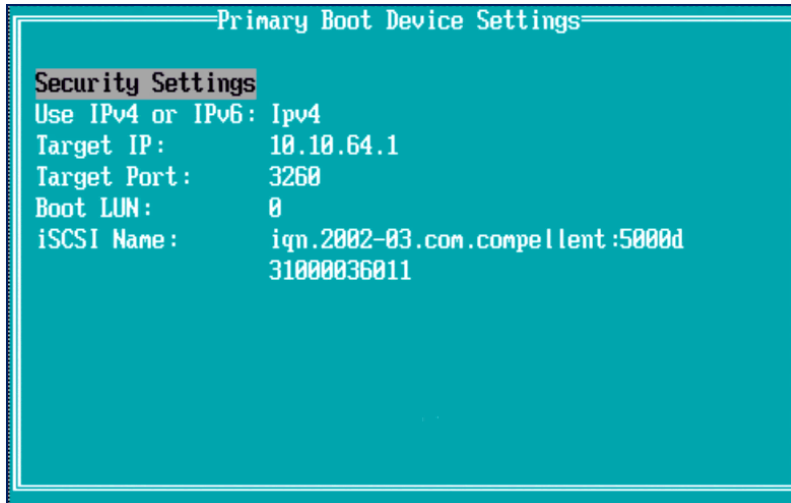


Figure 134. Primary Boot Device Settings

- 16) Back at the **iSCSI Boot Settings** menu select **Primary Boot Device Settings**.
- 17) Select **Use IPv4**.
- 18) Enter the appropriate IP address of the iSCSI controller for the Dell Compellent Storage Center.
- 19) Set the **Target Port = 3260**.
- 20) Set **Boot LUN = 0**.
- 21) Enter the **iSCSI Name** for the Dell Compellent iSCSI controller.
- 22) Hit **Escape** when completed.

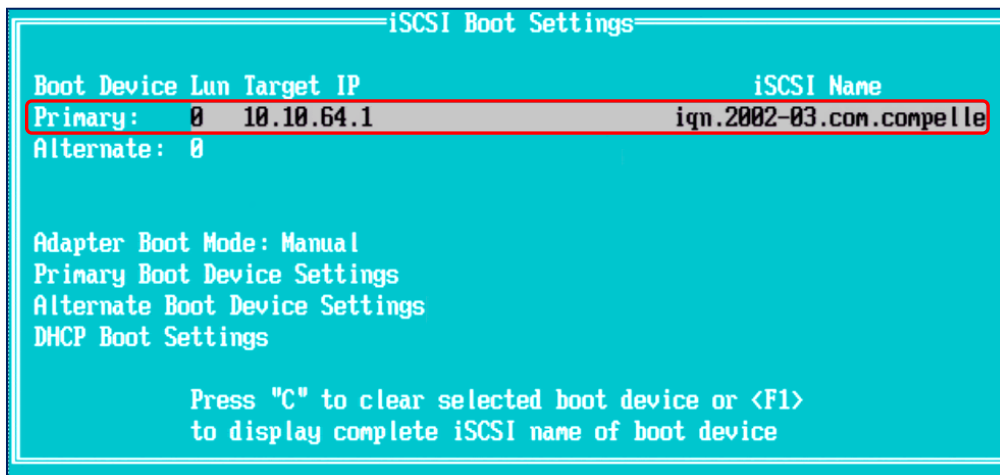


Figure 135. Primary Target LUN

- 23) On the **iSCSI Boot Settings** screen, select the **Primary** Boot Device.

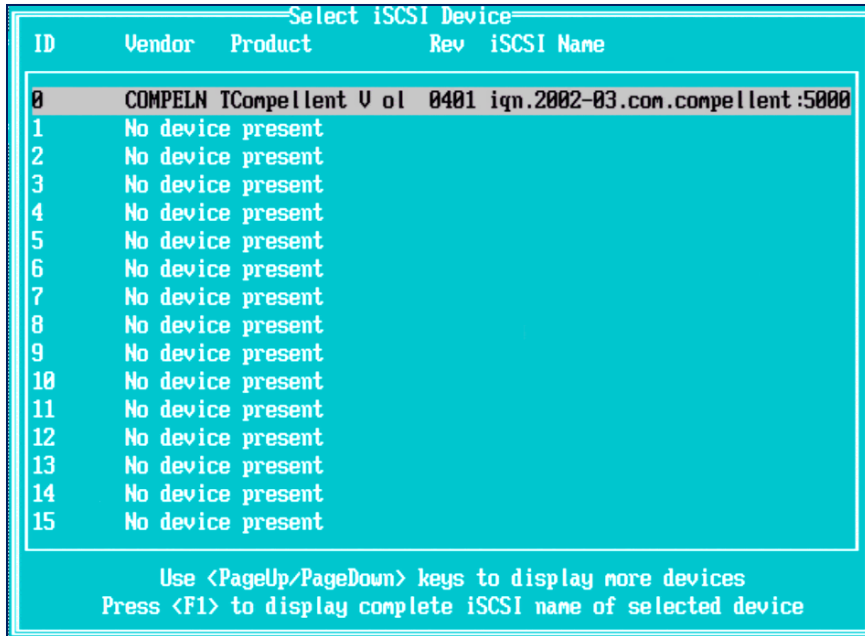


Figure 136. Select iSCSI Device

24) On the **Select iSCSI Device** screen, select **ID 0**.




Figure 137. Save Changes

25) Press **Escape** twice then select **Save changes**.

26) Reboot system and install the OS.

Note: Windows 2008 does not contain QLogic iSCSI drivers so these drivers will have to be loaded during the Windows 2008 setup from a mounted ISO file or a USB device.

27) If MPIO is desired, then install the MPIO feature, and repeat the above steps to enable a 2nd iSCSI path, and then configure MPIO settings. Please see Section 4 for more information on MPIO settings.

11 Storage Manager for SANs

The Dell Compellent VDS service works with Microsoft Storage Manager for SANs to simplify Storage management on Server 2008 Hyper-V host servers.

This procedures below describe how to install and configure both Microsoft Storage Manager for SANs and Dell Compellent’s VDS service. Both pieces of software must be installed prior to using Storage Manager for SANs to connect and configure Dell Compellent Storage.

11.1 Install Storage Manager for SANs

- 1) From a Windows 2008 server (full installation), login and launch **Server Manager**.
- 2) Select **Features** and then **Add Features**.

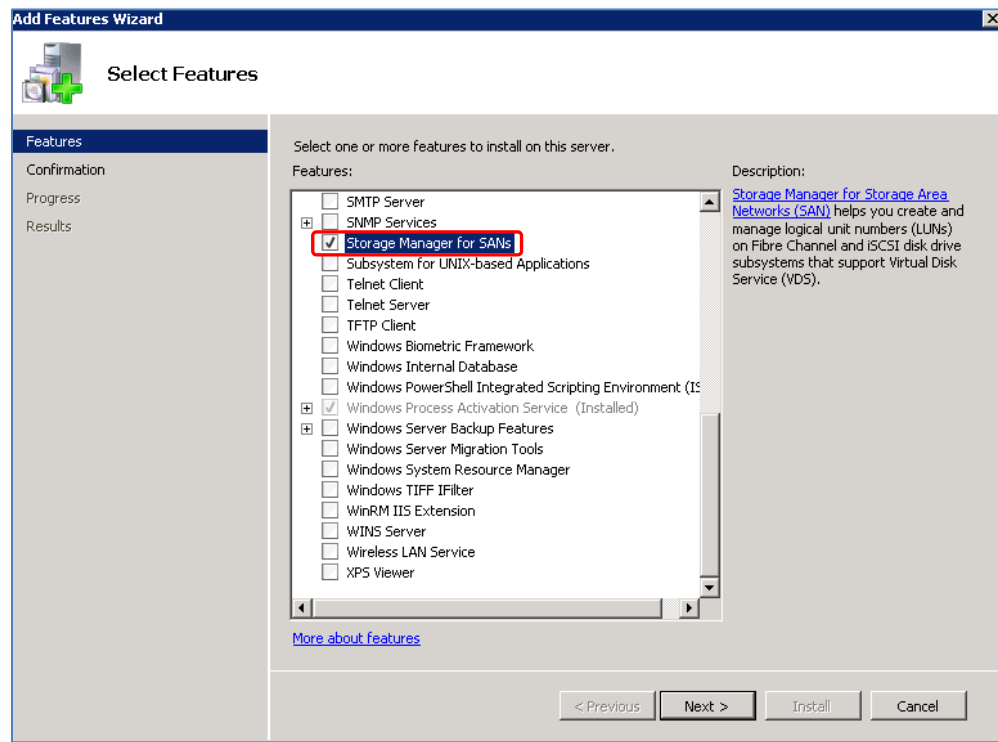


Figure 138. Storage Manager for SANs Feature

- 1) Check **Storage Manager for SANs** then click **Next**.
- 2) Click **Install** and allow the installation to finish.

11.2 Install Dell Compellent VDS (Virtual Disk Service)

- 1) Start the VDS setup by launching the **CompellentVDSManagerSetup.msi** installer file.
- 2) Click **Next** on the Welcome screen.
- 3) Review and accept the license agreement and click on **Next**.

- 4) Click **Next** on the **Destination folder** screen. Change the default directory if desired.
- 5) Click the **Install** button and allow the installation to finish.
- 6) When the install has completed, the **Launch VDS Manager Configuration** checkbox should be checked.
- 7) Click on the **Finish** Button to exit the install wizard and the configuration utility should start.

11.2.1 Configure Dell Compellent VDS Manager Service

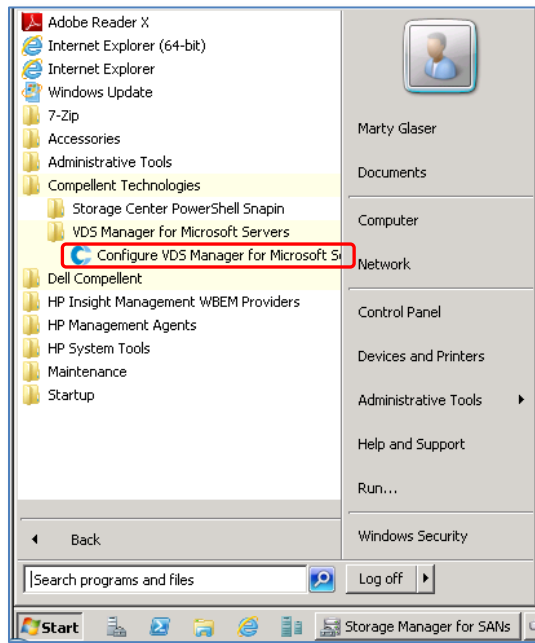


Figure 139. Configure VDS Manager

- 1) If the VDS configuration utility does not launch, then it can be run manually from **Start→All Programs→ Compellent Technologies→ VDS Manager for Microsoft Servers→Configure VDS Manager for Microsoft Servers.**
- 2) At the welcome screen, click on **Next**.

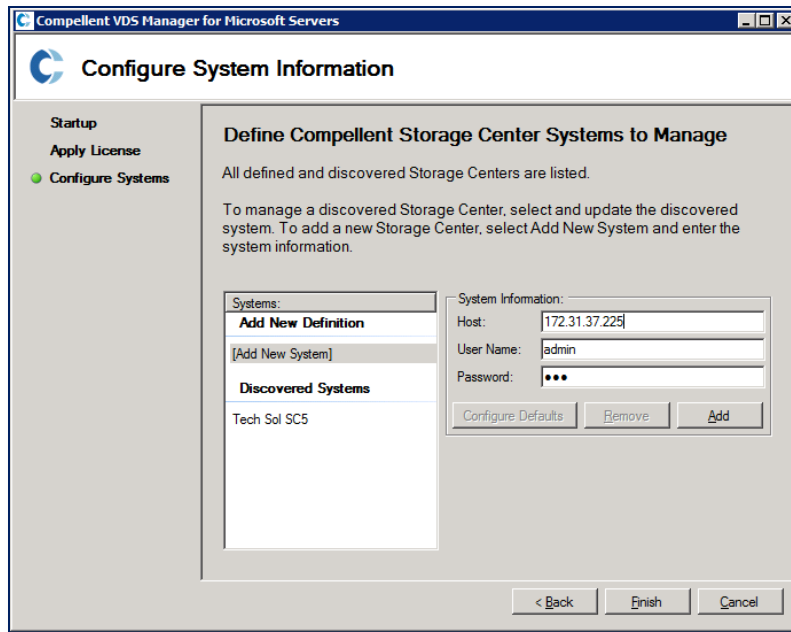


Figure 140. Define Storage Center

- 3) On the **Configure System Information** screen select **[Add New System]** and enter the appropriate login information for the Dell Compellent Storage Center then click the **Add** button.
- 4) Repeat step 3 if additional Dell Compellent Storage Center systems should be added.
- 5) When all the desired Storage Centers have been added, then click on **Finish**.

11.3 Storage Manager For SANs Example

In the example shown below, a LUN will be created using Storage Manager for SANs with Dell Compellent's VDS integration.

- 1) Once Dell Compellent VDS is installed and configured, launch **Storage Manager for SANs** from **Start→Administrative Tools**.

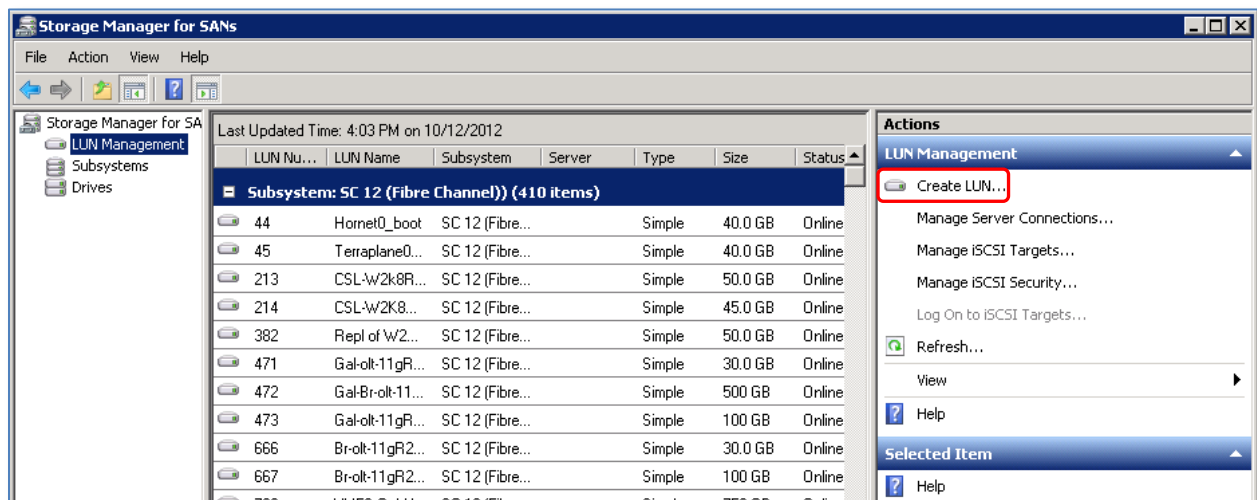


Figure 141. Create LUN with Storage Manager for SANs

- 1) To create a LUN, select **LUN Management** from the left pane and then select **Create LUN** from the **Actions** pane.
- 2) The **Provision Storage Wizard** screen should appear. Click **Next** to continue.

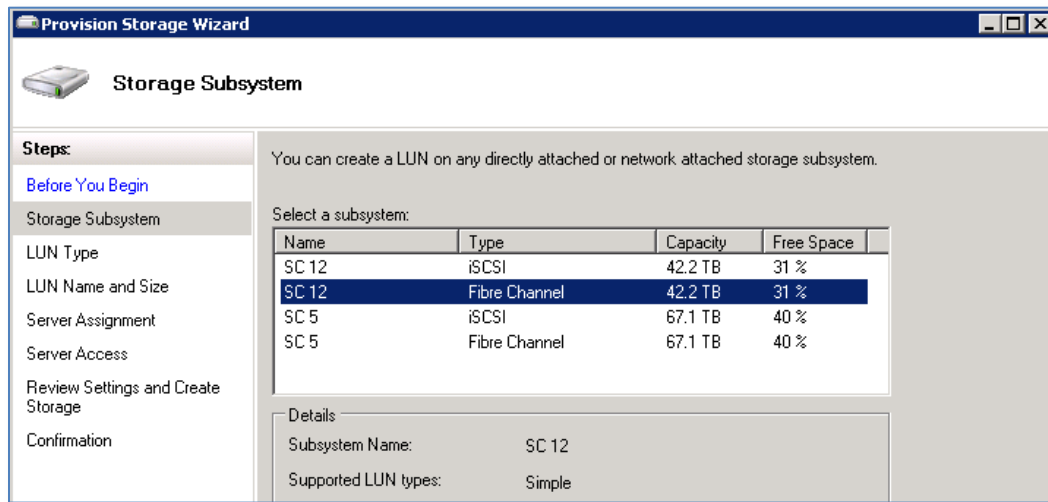


Figure 142. Select the Storage Subsystem

- 3) Select the desired Dell Compellent Storage Center to create the LUN on. In this example, Storage Center 12 (**SC 12 Fibre Channel**) is selected as shown in Figure 142. Click on **Next** to continue.

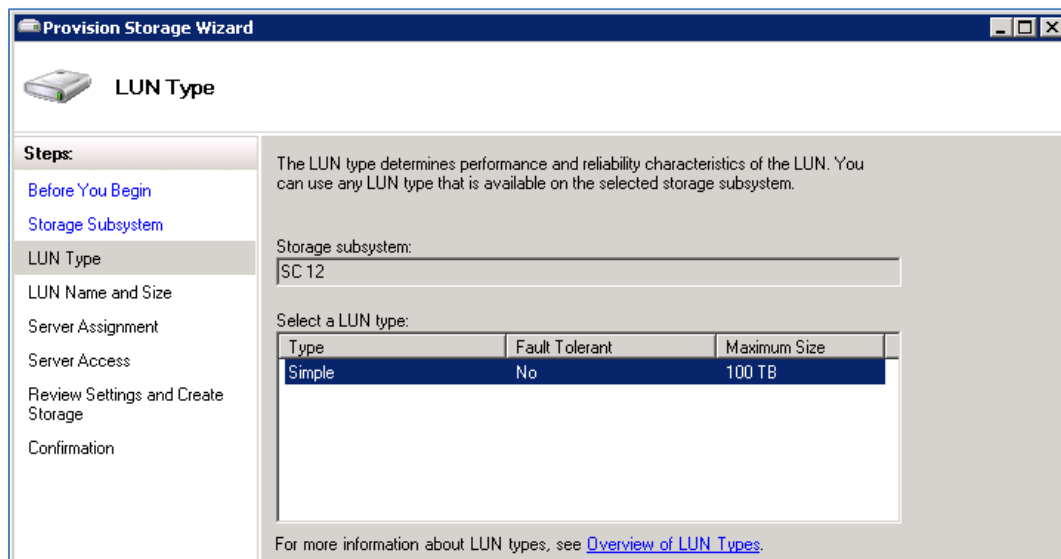


Figure 143. Select the LUN Type

- 4) On the **LUN Type** screen, highlight **Simple** then click on **Next**.

Note: "Simple" is the only option available on a Dell Compellent Storage Center.

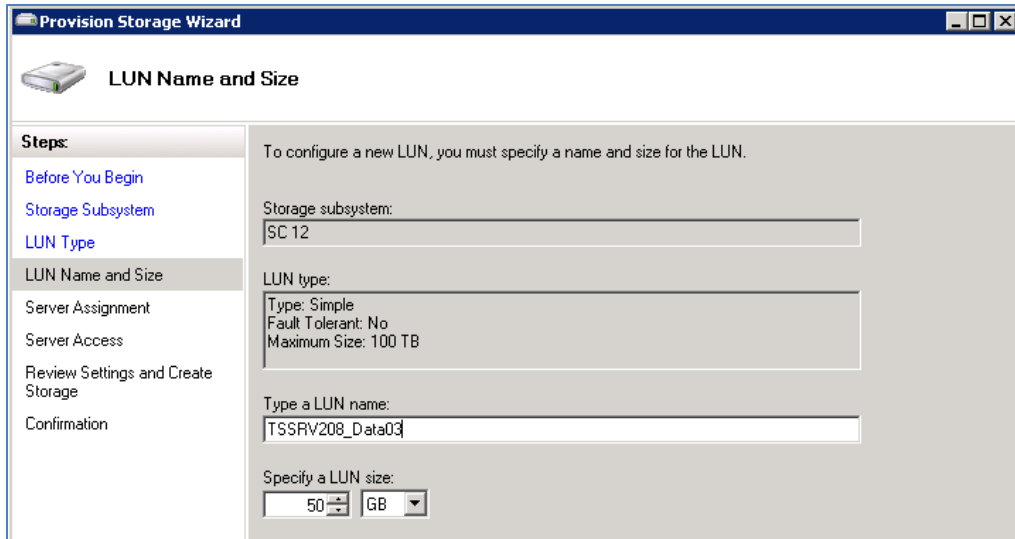


Figure 144. Define the LUN Name and Size

- 5) On the **LUN Name and Size** screen, provide a name for the new LUN and also specify the LUN size. Then click on **Next**.

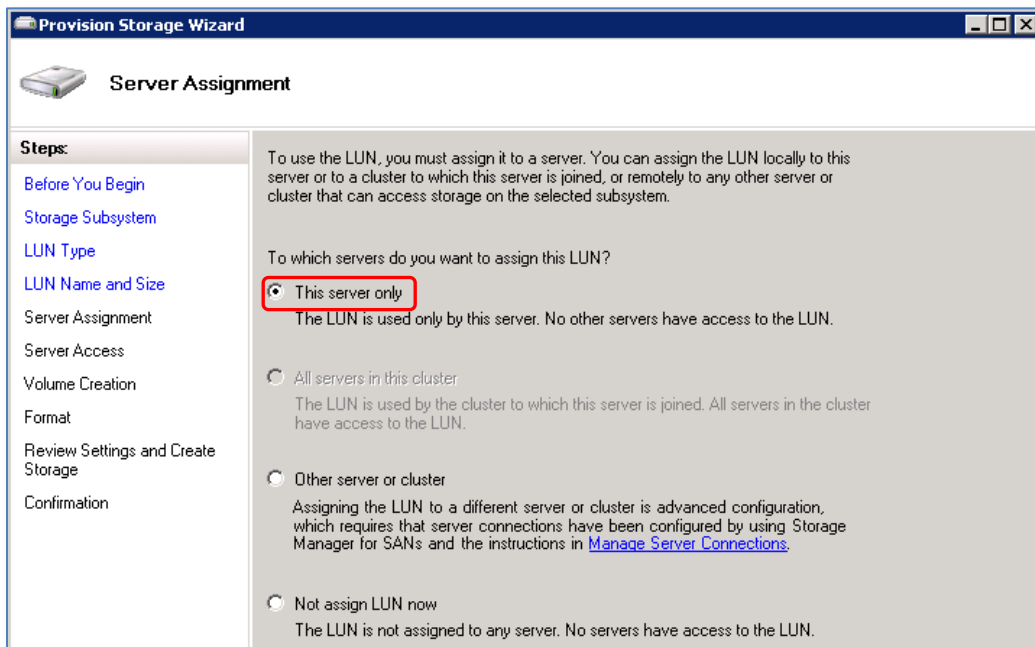


Figure 145. Select a Server Assignment

- 6) On the **Server Assignment** screen, select the appropriate option and then click on **Next**. In this example, **This server only** is selected.

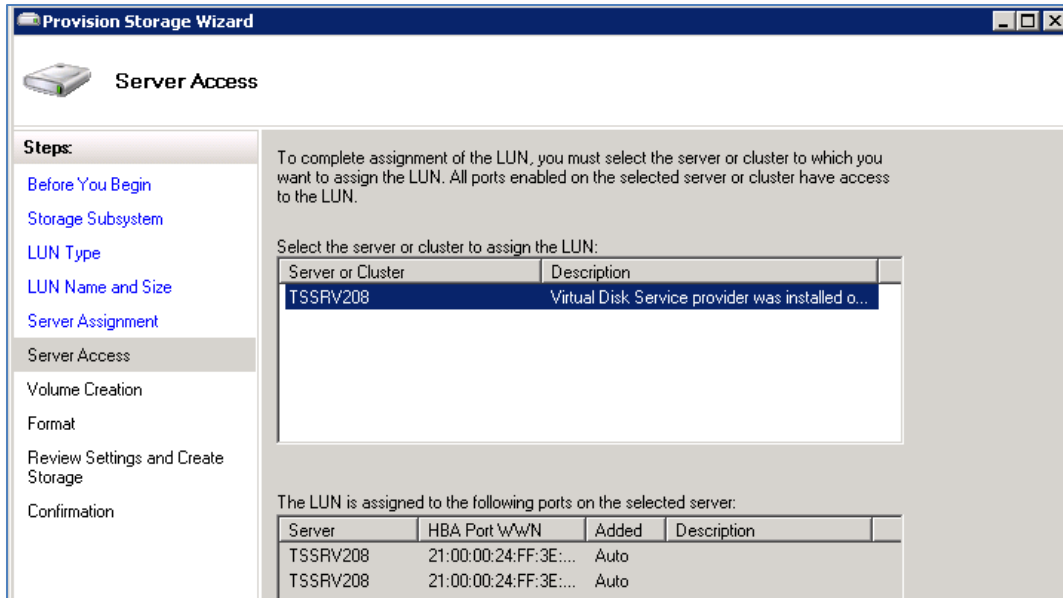


Figure 146. Define Server Access

- 7) On the **Server Access** screen, select the server (or server cluster) to assign the LUN to and then click **Next**.

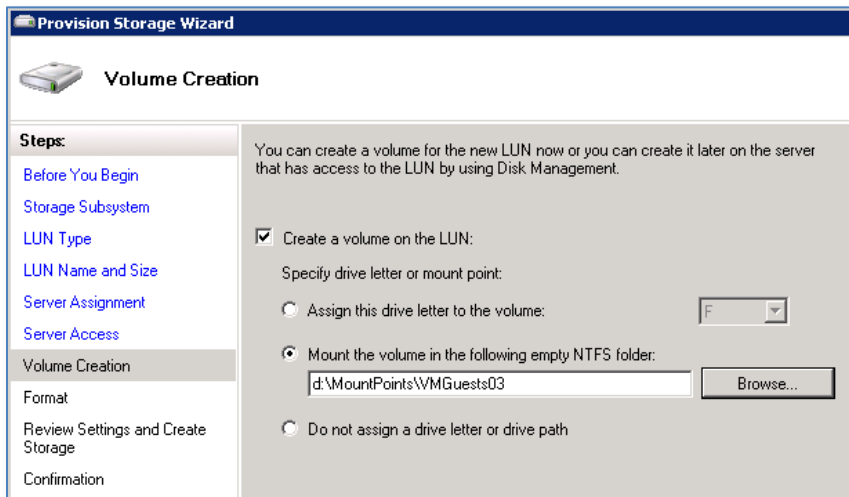


Figure 147. Create a Volume on a LUN

- 8) On the **Volume Creation** screen, assign a drive letter or define a mount point and then click **Next**. In this example, a mount point is specified.

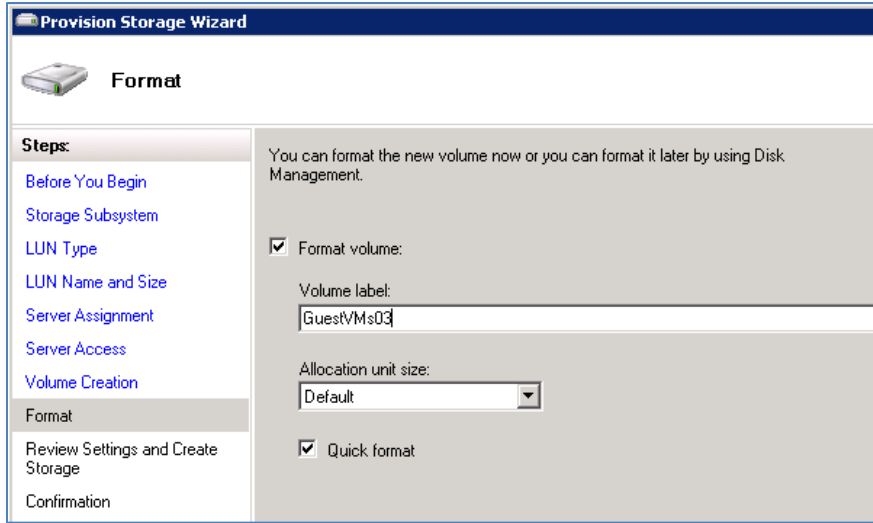


Figure 148. Provide a Volume Label, Allocation Size, and Format Option

- 9) On the **Format** screen, provide a **Volume label**, change the allocation size if desired, then click **Next**.

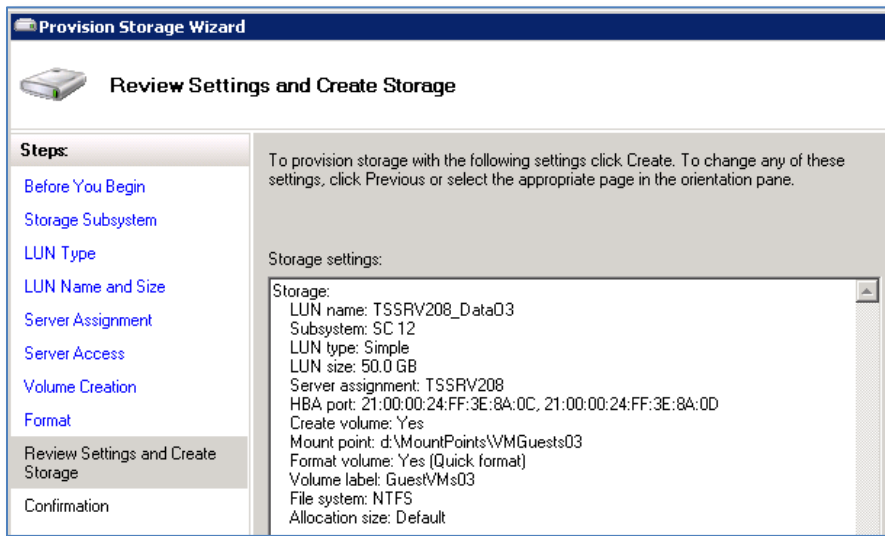


Figure 149. Review Settings and Create Storage

- 10) On the **Review Settings and Create Storage** screen, review the summary and then click **Create**.

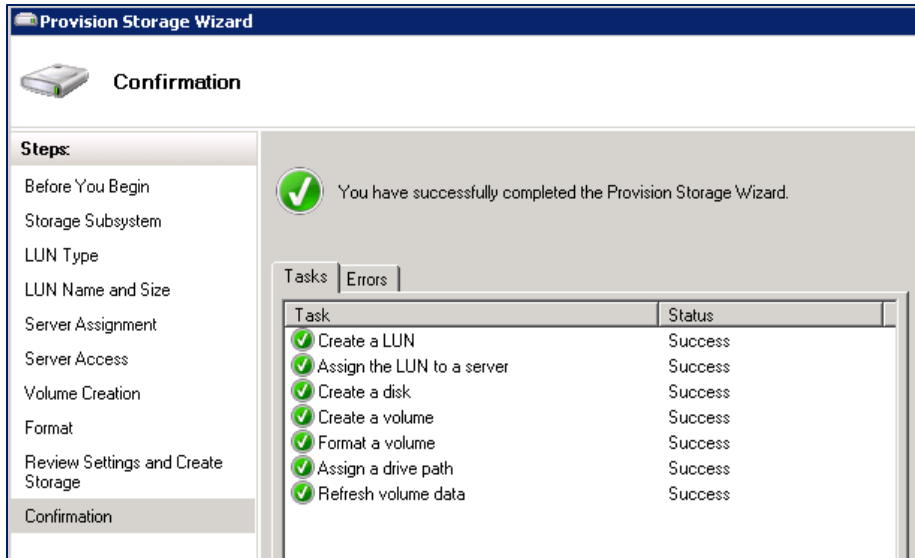


Figure 150. Confirmation Screen Shows Successful Creation of the New LUN

- 11) On the **Confirmation** screen, review the status column to verify that each task was successful, then click on **Close**.

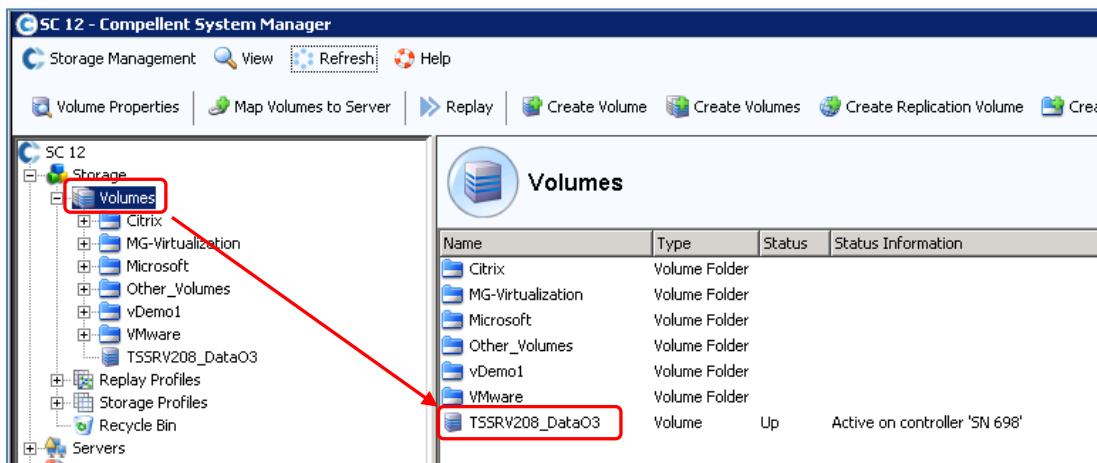


Figure 151. New Volume on Dell Compellent

- 12) When creating a new server volume with Storage Manager for SANs, the volume will be created at the root of the **Volumes** folder on Dell Compellent by default. If desired for management reasons, simply move the volume to a volume subfolder.
- 13) Ensure that a Storage Center Replay schedule is configured for the new volume, and enable replication of the new volume to another Storage Center if desired.

12 PowerShell Integration

Dell Compellent incorporates Windows PowerShell into its management feature set. This allows the ability to create, delete, map, and manage Dell Compellent volumes from PowerShell on a Hyper-V host server.

Administrators can create custom PowerShell scripts by using both the Dell Compellent PowerShell Cmdlet set and Hyper-V Cmdlet set together to quickly create volumes, map volumes to Hyper-V hosts, and create Hyper-V guest on those volumes.

For example, a PowerShell script can be used to automate the creation of new guest VMs from a sysprepped gold image. The script can map a View Volume containing the gold image to a Hyper-V host, re-scan the disks, add the View Volume as a new mount point, rename the virtual hard disk file, and create a new Hyper-V guest VM. When providing an input file, the script can be used to automate the deployment of many new guest VMs in a matter of minutes while maintaining space efficiency.

12.1 PowerShell Integration Example

Below is an example of a PowerShell Script that:

- Inputs a VM guest name TXT file
- Uses a gold image Storage Center Replay as the source for View Volumes
- Indicates the Hyper-V host to place the new guest VM servers on
- Indicates where to place the mount points on the Hyper-V host server

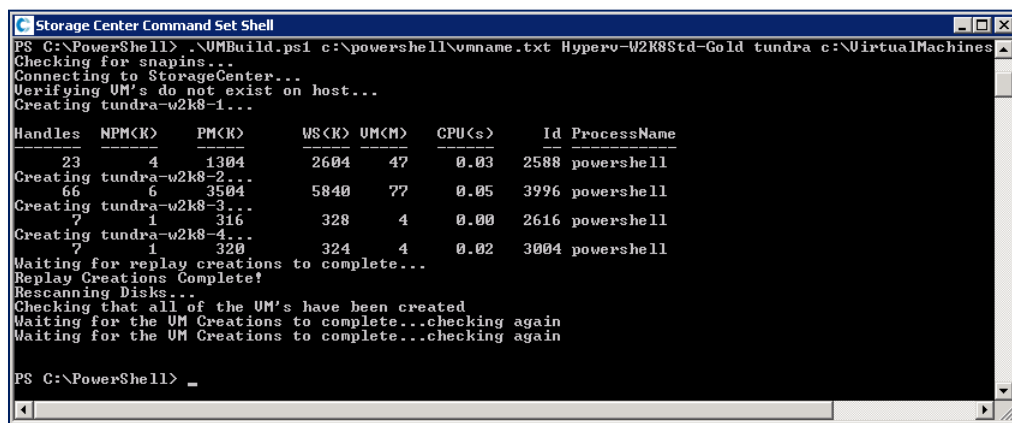


Figure 152. Storage Center Command Set Shell

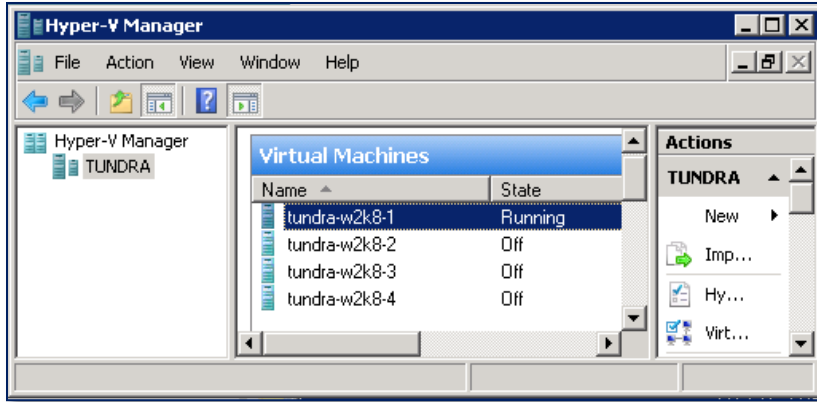


Figure 153. New guest VMs Created from a PowerShell Script

In Figures 152 and 153, three new View Volumes were created from a gold image Replay and used to provision three new guest VMs on the Hyper-V host server.

In addition, Server 2012 offers over 150 new built-in Cmdlets to perform Hyper-V management tasks.

13 Dynamic Memory

Dynamic Memory was introduced with Server 2008 R2 SP1 Hyper-V and enhanced with Server 2012. An overview and some best practices are offered below.

13.1 Dynamic Memory Configuration

With Windows Server 2008 R2 SP1 and Server 2012, the physical RAM on a Hyper-V host can be dynamically allocated to (and reclaimed from) Hyper-V guest VMs. In earlier versions of Hyper-V, a fixed amount of RAM had to be assigned to each guest VM. This usually resulted in some inefficient use of RAM on the guests and the host due to having to set the RAM on each guest according to the peak RAM that guest VM may need.

With Dynamic RAM, an administrator can specify a minimum and maximum amount of RAM available to each guest VM. If a guest needs more RAM during peak usage times, it can acquire it from the pool of free RAM on the host server (up to the guest VM's max RAM setting or until the host server's pool runs out, whichever comes first).

When the guest VM server no longer needs the extra RAM, it can be reclaimed by the host server so it can be allocated to other guest VMs.

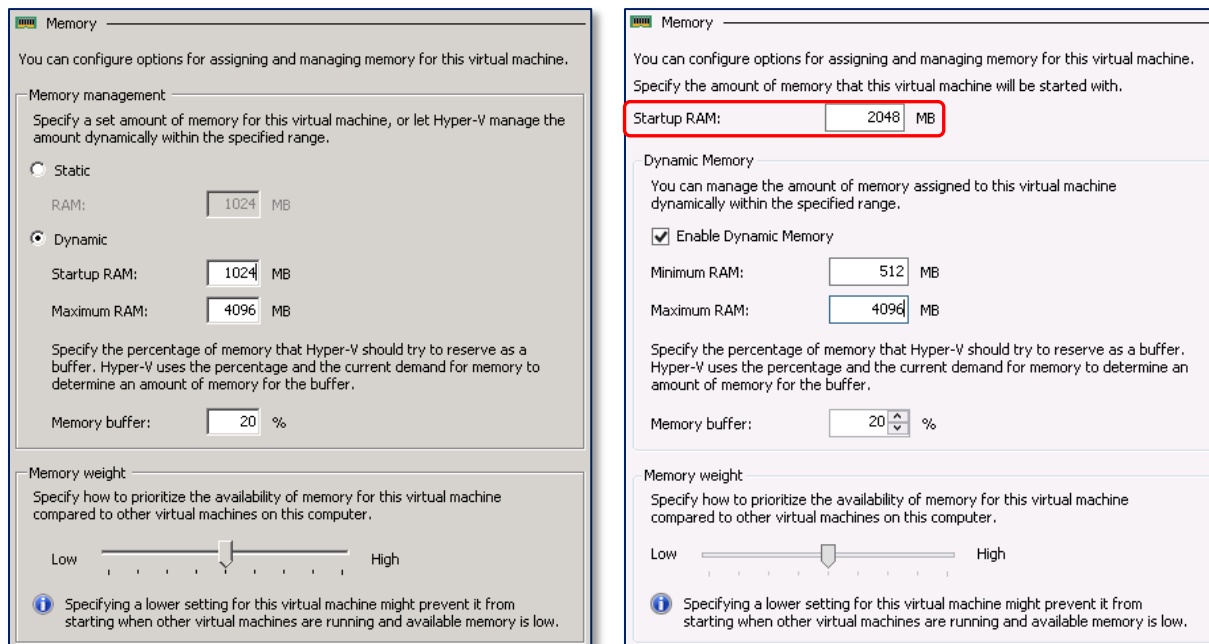


Figure 154. Comparing Server 2008 R2 SP1 and Server 2012 Hyper-V Dynamic Memory

When comparing the dynamic memory settings for Server 2008 R2 SP1 Hyper-V (left) and Server 2012 Hyper-V (right) in Figure 154, there are two main differences:

- Server 2012 (right) adds the ability to specify more startup RAM than the minimum

RAM. Being able to set a startup RAM value that is higher than the minimum RAM ensures that the guest VM is not starved for RAM when initially booting with the minimum RAM setting. After the server boots, RAM that is not needed by the guest VM for the boot process can be returned to the host server's RAM pool to be used by other guest VMs.

- With Server 2012, a guest VM no longer needs to be powered off to change the minimum and maximum RAM settings. However, the amount of startup RAM can be changed only when the guest VM is powered off.

13.2 Memory Settings for Hyper-V Guests Running Citrix

A Hyper-V guest VM running Citrix loads a user profile for each authenticated user. If the memory footprint for each profile is 256 MB of RAM, then 35 concurrent Citrix users would require about 9 GB of RAM (256MB x 35 users).

When the users end their sessions, the 9 GB of RAM is freed up and returned to the pool on the host server to be used by other guest VMs.

In deciding how to configure dynamic RAM settings for a guest VM running Citrix in the above example, given a Server 2012 host, an administrator might specify a startup RAM setting of 2048 MB for optimal boot performance, a minimum RAM setting of 1024 MB, and a maximum RAM setting of 10240 MB, and tune as necessary.

13.3 Memory Settings for Hyper-V Guests running SQL

Some applications such as SQL may by design consume the free RAM available to a guest VM simply because it is there for the taking (regardless of whether the application actually needs the RAM) and not return the RAM to the host to be used elsewhere.

For example, if it is known that a particular guest VM running SQL server will never need more than 2 GB of RAM to run optimally, the guest should not be assigned more than 2 GB RAM. If an administrator sets a maximum RAM setting of 4 GB on that guest VM, SQL may consume RAM beyond the 2GB it needs, simply because the RAM is there for the taking, with without realizing any measurable increase in performance for consuming the extra RAM.

If the guest is running on Server 2012 Hyper-V, if a change (such as increased user count) requires the guest VM's maximum available RAM be increased from 2 GB to 4 GB, an administrator can quickly allocate this RAM without having to power down the guest.

Where possible, administrators should adjust application settings to ensure optimal use of RAM, with the goal that they only use the RAM they need, and release RAM that is no longer needed back to the guest so that the host server can assign it elsewhere.

In addition, memory weight and buffer settings allow for some additional tuning of guests to give priority to some guests over others if the amount of free RAM on the host ever does run low.

Table 5. Comparing Fixed and Dynamic Memory with Hyper-V

Without Dynamic Memory		With Dynamic Memory		
Total Host Server RAM	16 GB	Total Host Server RAM	16 GB	
RAM Reserved for Host	2 GB	RAM Reserved for Host	2 GB	
RAM available for Guests	14 GB	RAM available for Guests	14 GB	
VM1	1 GB	VM1	.5 GB (min)	1 GB (max)
VM2	1 GB	VM2	.5 GB (min)	1 GB (max)
VM3	2 GB	VM3	.5 GB (min)	2 GB (max)
VM4	2 GB	VM4	.5 GB (min)	2 GB (max)
VM5	8 GB	VM5	.5 GB (min)	8 GB (max)
----	----	VM6	.5 GB (min)	2 GB (max)
----	----	VM7	.5 GB (min)	1 GB (max)
----	----	VM8	.5 GB (min)	1 GB (max)
----	----	VM9	.5 GB (min)	1 GB (max)
----	----	VM10	.5 GB (min)	1 GB (max)
Total	14 GB	Total	5 GB (min)	20 GB (max)

In the example shown in Table 5, a Hyper-V host server has 16 GB of RAM. 2 GB is reserved for the host, and 14 GB is available for guests. Without dynamic memory, the host is able to accommodate fewer guests (only 5 guests as shown in this example).

With dynamic memory, more guest VMs can be provisioned (5 more in this example, for a total of 10 guests). Note that the host server itself is over-provisioned by 6 GB.

Note: While dynamic memory with Hyper-V allows for better RAM utilization, it also involves some risk due to being able to over-provision the RAM on the host. Before implementing dynamic memory, administrators should have a good understanding of the behavior of their guest VMs to avoid exhausting the available RAM pool and causing an interruption to service. The information in Table 5 is provided as an example only. Individual results will vary depending on factors that are unique to each environment.

14 Virtual Fiber Channel

One of the new features introduced with Server 2012 Hyper-V is the ability for guest VMs to use virtual fiber channel adapters to connect to SAN volumes directly.

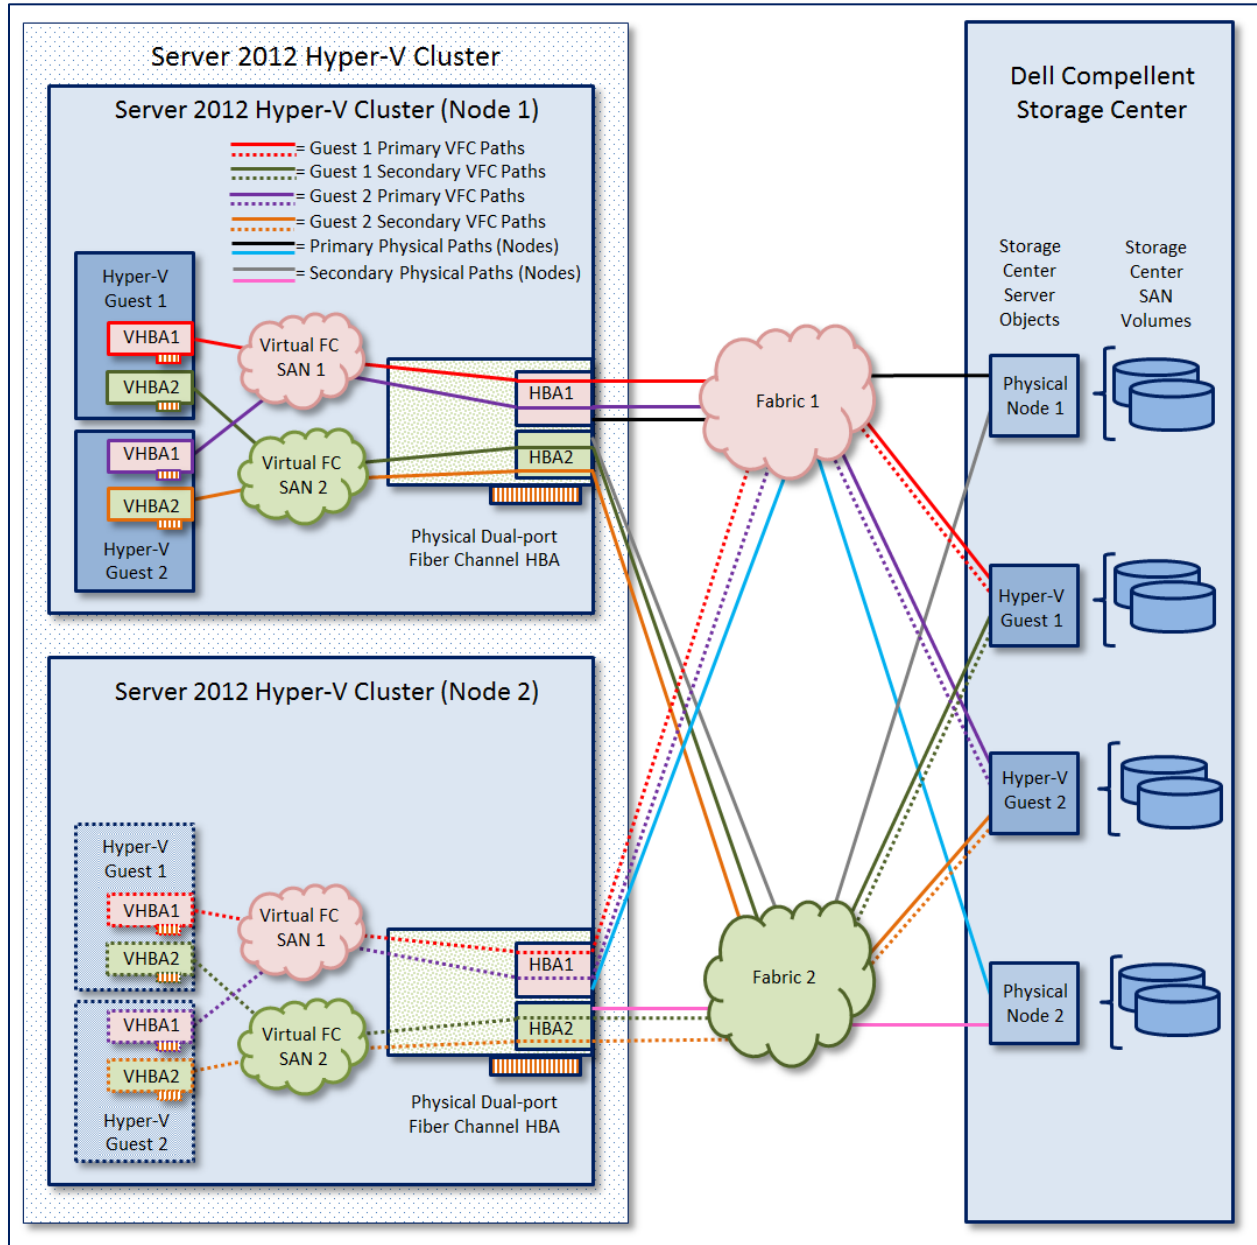


Figure 155. Virtual Fiber Channel for Server 2012 Hyper-V Guests

As shown in Figure 155, virtual fiber channel adapters allow administrators to present Dell Compellent storage volumes directly to Hyper-V guest VMs running Server 2008 R2 or Server 2012.

Note: virtual fiber channel requires that the physical fiber channel HBAs in the Hyper-V hosts support N-Port ID Virtualization (NPIV).

Server 2012 Hyper-V virtual fiber channel allows a physical fiber channel HBA that is NPIV-capable to pass world-wide-names (WWNs) from the guest VMs to the fabric so that guest VM server objects can be created on a Dell Compellent Storage Center and mapped directly to SAN volumes.

Figure 155 also illustrates that MPIO is fully supported with virtual fiber channel on guest VMs, and that virtual fiber channel is supported on physical Hyper-V clusters to allow guest VMs with virtual fiber channel adapters to live migrate to other nodes in the cluster. In Figure 155, Node 2 shows that placeholder paths (dashed lines) are in place, and these paths become active when the guest VM is live-migrated to that node.

Table 6. Guest VM Clustering Options

Type of Hyper-V Node	Supported Topologies for Guest VM Clustering	
	iSCSI	Virtual Fiber Channel
Server 2008 R2	✓	
Server 2012	✓	✓

As shown in Table 6, with server 2008 R2 Hyper-V, guest VMs can be clustered by using direct-attached iSCSI SAN volumes. With Server 2012, guest VM clustering support has been extended to include virtual fiber channel. This opens up new possibilities and scenarios for configuring high-availability with guest VMs when fiber channel is the preferred topology.

14.1 Configure a Hyper-V Guest to use Virtual Fiber Channel

In the following example, a Server 2008 R2 guest VM (Server 2012 is very similar) on a two-node cluster will be configured to use virtual fiber channel, and storage will be provisioned on a Dell Compellent Storage Center and presented to the guest VM, incorporating MPIO.

The below example assumes the following environment is already in place:

- Two physical Server 2012 Hyper-V hosts in a clustered configuration.
- The physical host servers have physical fiber channel HBAs that support NPIV. In this example, the host servers each have a dual-port QLogic QLE2562 HBA card.
- The host servers have the MPIO feature installed and configured (round robin).
- A dual fiber fabric that supports NPIV is available to provide full data path redundancy.
- A Server 2008 R2 or Server 2012 guest VM is installed on a node in the cluster.
- The guest VM is booting from a virtual hard disk (VHDX) file that is presented to the guest VM as a virtual IDE drive (the boot drive must be a virtual IDE device).
- A Dell Compellent Storage Center with available storage capacity is available to the cluster and the guest VM.

CAUTION: There is a fabric configuration dependency to be aware of when planning to deploy virtual fiber channel for guest VMs.

- It is important to ensure that your physical fiber channel fabrics use only “World Wide Name” (WWN) zoning for all your physical Hyper-V nodes and Storage Centers (not “Switch Port Name” zoning) before implementing virtual fiber channel.
- Attempting to boot a guest VM with a virtual fiber channel HBA may result in one or more physical nodes in that cluster crashing (causing a service outage) if the fabric zoning method is Switch Port Name instead of WWN.
- To ensure compatibility, please verify with your fabric administrator that WWN zoning is configured for your Hyper-V nodes and Storage Centers before completing the steps below.

14.1.1 Create Virtual SANs on the Hyper-V Nodes

The first step is to create virtual fiber-channel (FC) SANs using Hyper-V Manager. Two virtual FC SANs will be created since both nodes have dual fiber channel HBAs associated with a dual fabric, as shown in the example in Figure 155. Once the virtual FC SANs are in place on the host servers, any guest (that supports virtual fiber channel) can be configured to use these virtual FC SANs to connect to Dell Compellent volumes directly.

- 1) On the first Server 2012 Hyper-V node in the cluster, launch Hyper-V Manager and under the **Actions** screen, select **Virtual SAN Manager**.

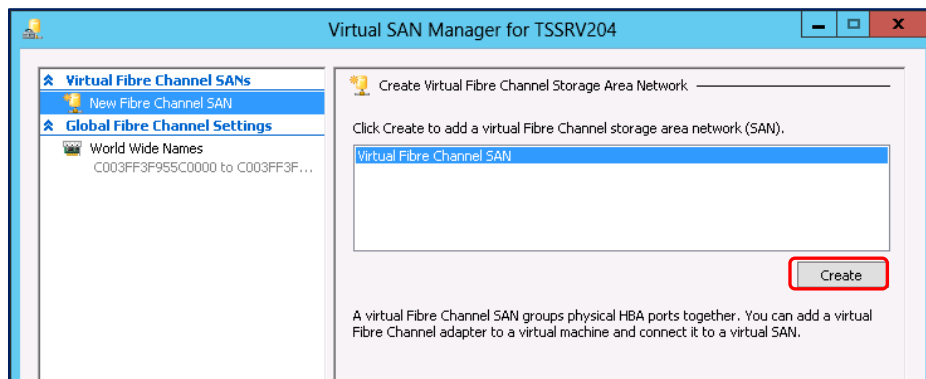


Figure 156. Virtual SAN Manager

- 2) Because both the server nodes in this example are configured with dual fiber channel HBAs connected to dual fabrics with MPIO configured, the guest servers will also be configured to use dual (virtual) fiber channel HBAs and MPIO.
- 3) With **New Fibre Channel SAN** highlighted, click on the **Create** button.

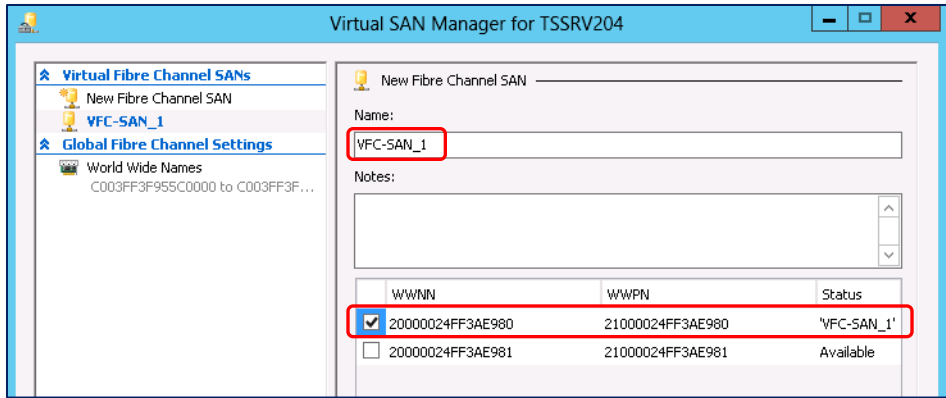


Figure 157. Create a Virtual SAN

- 4) Provide a descriptive name for the virtual SAN. In this example, **VFC-SAN_1** is used (Virtual Fiber Channel SAN 1).
- 5) Check the box in front of the **WWNN/WWPN** that is associated with the physical host's **HBA1** as shown in Figure 155. Then click on **Apply**.

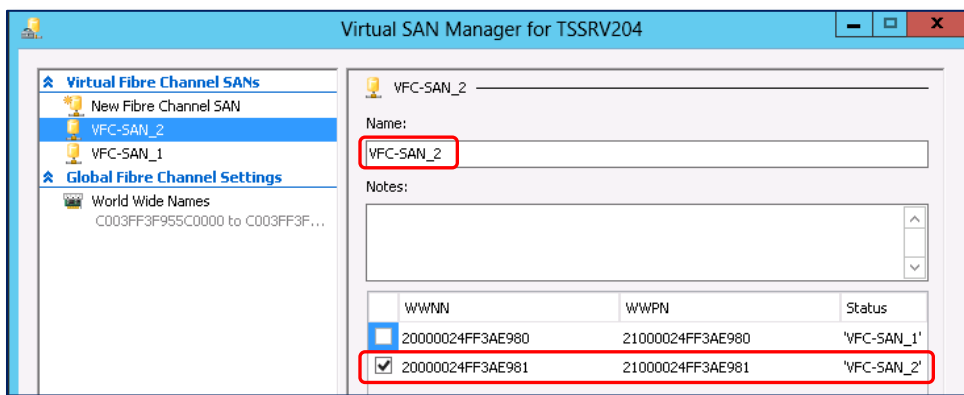


Figure 158. Two Virtual SANs Created for the Guest VM

- 6) Highlight **New Fibre Channel SAN** a second time, and repeat steps 3 - 5. The second virtual SAN in this example was named **VFC-SAN_2** and associated with the **WWNN/WWPN** for the physical host's **HBA2**.
- 7) Click on **OK** to close out of the **Virtual SAN Manager** window.
- 8) Repeat steps 1 – 7 above on each additional node in the cluster (there are two nodes in this example).

Note: make sure that the Virtual Fiber Channel SAN names (**VFC-SAN_1** and **VFC-SAN_2** in this example) are exactly the same on all nodes of the cluster.

14.1.2 Prepare a Guest VM to Use Virtual Fiber Channel Adapters

Now that the virtual FC SANs are in place on the nodes of the cluster, the guest VM has to be configured to use virtual fiber channel adapters and MPIO.

- 1) On the desired guest VM (Server 2008 R2 in this example), go to **Start→Administrative Tools→Server Manager** and install **Multipath I/O** under **Features**. If the guest will also be clustered, install the **Failover Clustering** feature.
- 2) Verify that the latest integration services have been installed on the guest VM. For more information on installing integration services, please refer to Section 9.2.2.
- 3) Power the guest VM off.

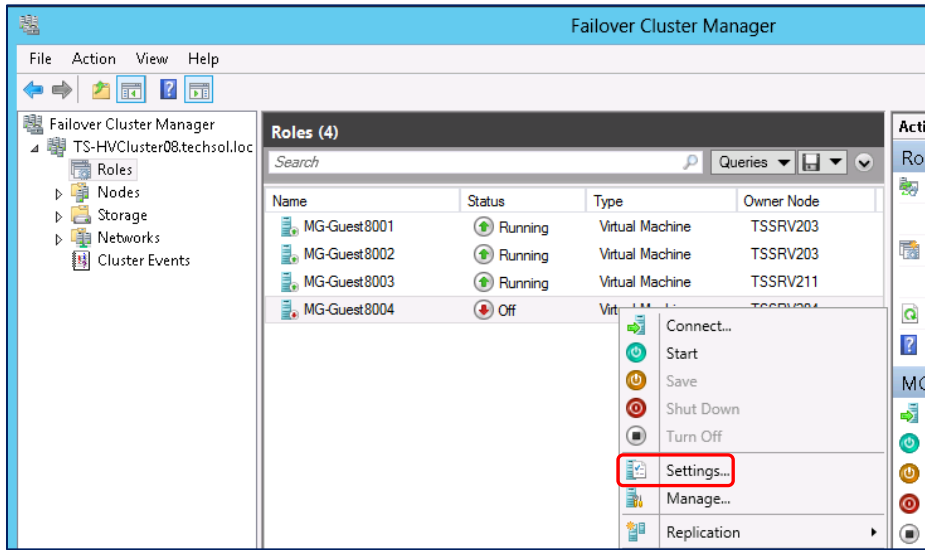


Figure 159. Edit Guest VM Settings

- 4) From a node in the cluster, launch **Failover Cluster Manager** and right-click on the desired guest VM and choose **Settings**. In this example, the guest server **MG-Guest8004** will be configured to use virtual fiber channel.

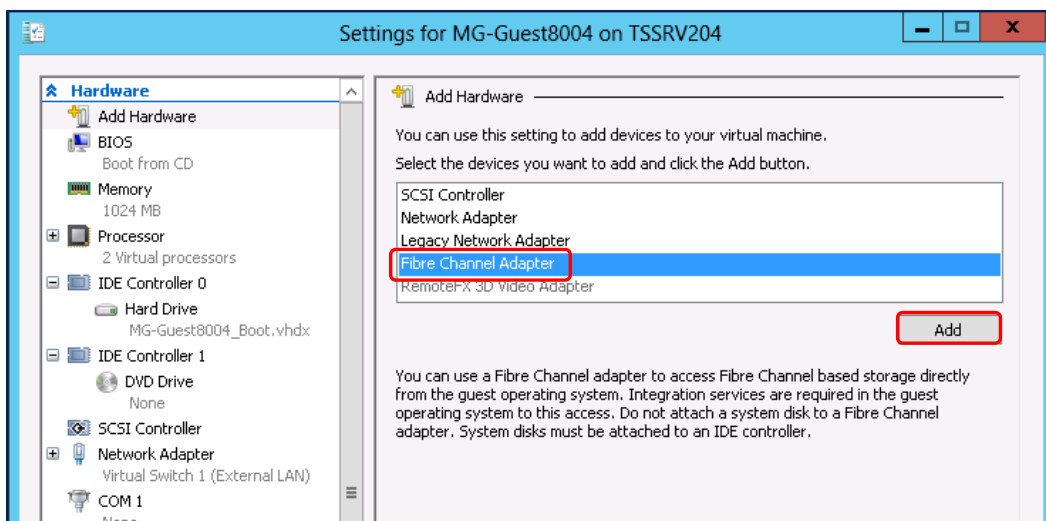


Figure 160. Add a Fiber Channel Adapter

- 5) Select **Add Hardware** at the top of the **Hardware** column, and on the right, select **Fibre Channel Adapter** and then click on **Add**.

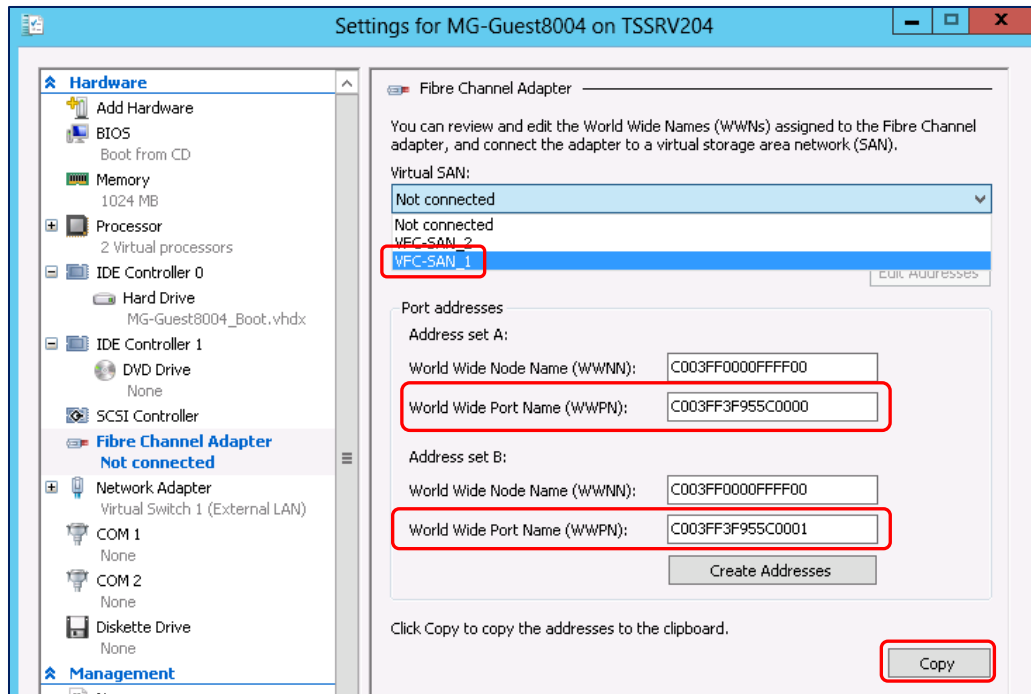


Figure 161. Add the First Virtual Fiber Channel Adapter to the Guest VM

- 6) From the drop-down list under **Virtual SAN**, select the first virtual SAN (**VFC-SAN_1** in this example) and then click on **Apply**.
- 7) Make note of the two WWPNs listed for this virtual HBA. These are the WWPNs that will be visible from this virtual HBA to the fabric and to the Dell Compellent Storage Center for mapping purposes in the steps that follow below.
- 8) The WWNN and WWPN addresses that are suggested by Failover Cluster Manager for the guest VM can be changed to addresses of your choice if desired.
- 9) By default, Failover Cluster Manager will not vary the World Wide Node Name (WWNN) – it stays the same for all virtual HBAs for all guest VMs on the physical cluster.
 - a. If the guest VMs using virtual fiber channel adapters will be clustered, then it is important to change the WWNN for each guest to make it unique. To do this, click on the **Edit Addresses** button.

Note: If the WWNN is not changed manually so it is unique for each guest VM, the **SCSI-3 Persistent Reservations** cluster validation test will fail, preventing the guest VMs from being clustered.

- b. If the guest VM using virtual fiber channel adapters will not be clustered with another guest VM, then the same WWNN can be used for all guests, but the best practices recommendation is to change the WWNNs to ensure that they are unique for each guest VM so they don't have to be changed later if clustering them becomes necessary.

- 10) Once satisfied with the WWNNs and WWPNNs, for convenience, click on the **Copy** button to copy these addresses to the clipboard, and paste them into a document for later reference. This will make zoning steps easier in the steps that follow below.

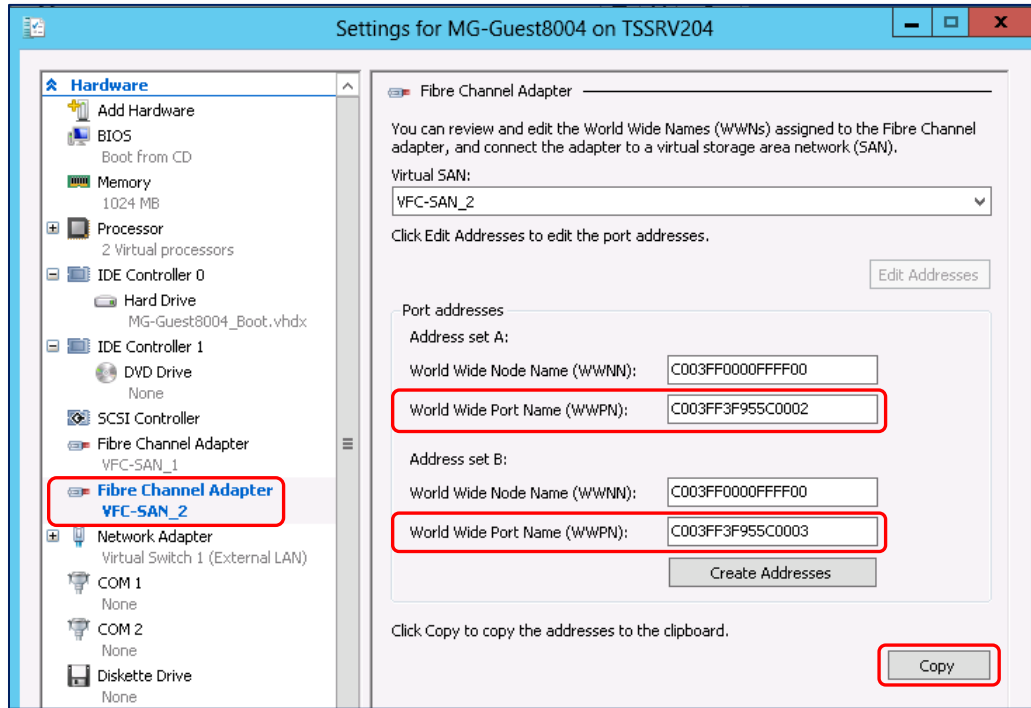


Figure 162. Add a Second Virtual Fiber Channel Adapter to the Guest VM

- 11) Add a second fiber channel adapter to the guest VM by repeating steps 5 – 10 above. Choose the second Virtual SAN from the drop down list (**VFC-SAN_2** in this example).

Note: Ensure that a unique WWNN is defined if this guest will be clustered with another guest VM.

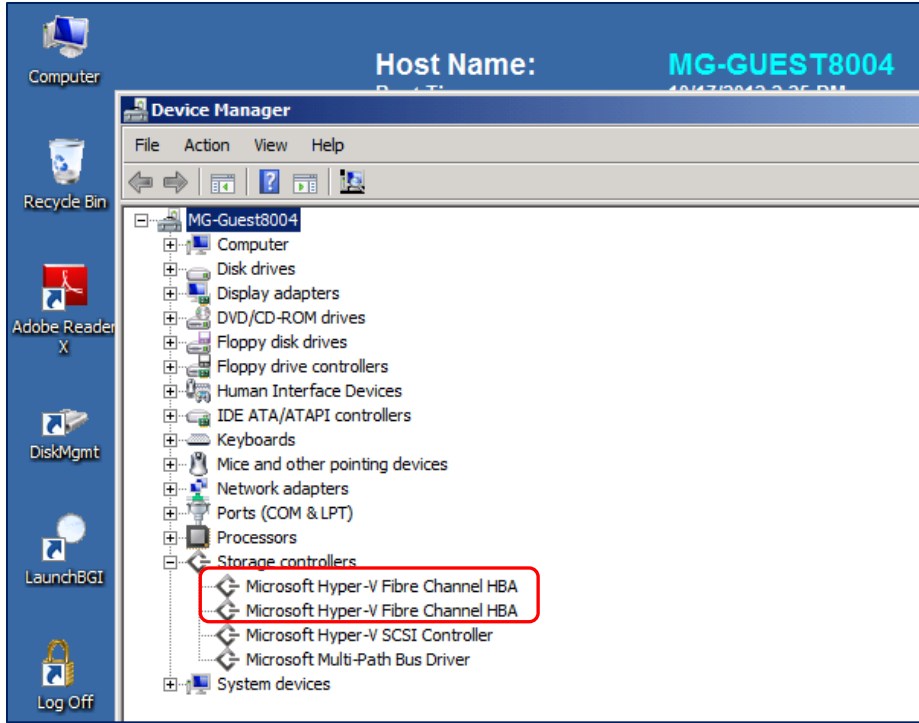


Figure 163. Guest VM with Microsoft Hyper-V Fiber Channel HBAs in Device Manager

12) Power on the guest VM and log in. Go to **Device Manager** and verify that the guest VM can see two **Microsoft Hyper-V Fiber Channel HBAs** as shown above.

Note: For Server 2008 R2 guests running on Server 2012 Hyper-V, the latest integration services must be installed on the guest VM before the virtual HBAs will display correctly in Device Manager. See Section 9.2.2 for more information on integration services.

14.1.3 Configure Zones for Fiber Fabrics

At this point, the guest VM's active WWPNs should be visible to the physical fiber fabrics. If required for your fabric topology, complete any required zoning steps for your fabrics, referring to Tables 7 and 8 below. If no zoning steps are required, please skip to section 14.1.14.

Table 7. Active WWPNs when MG-Guest8004 is on Node 1

Virtual HBA	Address Set	WWPN	Fabric Zones
Virtual FC Adapter 1 on VFC-SAN_1	A	C003FF3F955C0000	Fabric 1: Zone Members for "MG-Guest8004_Fabric-1" <input checked="" type="checkbox"/> Alias for Compellent Front-end SAN ports (for Fabric 1) <input checked="" type="checkbox"/> C003FF3F955C0000 <input checked="" type="checkbox"/> C003FF3F955C0001
	B	C003FF3F955C0001	
Virtual FC Adapter 2 on VFC-SAN_2	A	C003FF3F955C0002	Fabric 2: Zone Members for "MG-Guest8004_Fabric-2" <input checked="" type="checkbox"/> Alias for Compellent Front-end SAN ports (for Fabric 2) <input checked="" type="checkbox"/> C003FF3F955C0002 <input checked="" type="checkbox"/> C003FF3F955C0003
	B	C003FF3F955C0003	

Table 7 shows that two WWPNs (the “A” set) for MG-Guest8004 are active (highlighted in yellow) when the guest VM is on Node 1 of the Hyper-V cluster. Note that the other two WWPNs (the “B” set) are off line and not visible to the fabric when the guest VM is on Node 1.

Table 8. Active WWPNs After MG-Guest8004 is Live Migrated to Node 2

Virtual HBA	Address Set	WWPN	Fabric Zones
Virtual FC Adapter 1 on VFC-SAN_1	A	C003FF3F955C0000	Fabric 1: Zone Members for “MG-Guest8004_Fabric-1” <input checked="" type="checkbox"/> Alias for Compellent Front-end SAN ports (for Fabric 1) <input checked="" type="checkbox"/> C003FF3F955C0000 <input checked="" type="checkbox"/> C003FF3F955C0001
	B	C003FF3F955C0001	
Virtual FC Adapter 2 on VFC-SAN_2	A	C003FF3F955C0002	
	B	C003FF3F955C0003	

When the guest VM is live-migrated to another Node in the Hyper-V cluster, the guest VM will flip from the “A” WWPN pair to the “B” WWPN pair as part of the live migration process, and the “B” pair (as highlighted in yellow) becomes active on the fabric, as shown in Table 8 above.

Each time the guest VM is live migrated to another node in the cluster, it will alternate back and forth between the “A” and “B” WWPN pairs so that only one pair is active at a time. This is very important to understand because when creating zones for the physical fabrics 1 and 2 (as shown in the example in Figure 155) , all four WWPNs must be included. Having all 4 WWPNs included is necessary in order for order for live migration to work.

Because only two of the guest’s four WWPNs are active at any given time, only those two WWPNs will be visible on the fabric. Adding the inactive pair (the “down” WWPNs) to your fabric zones has to be done manually.

- Type the WWPNs for the “down” HBAs in manually when completing zoning steps.
- To avoid having to manually type in the WWPNs, copy and paste them from the document created in step 10 in the previous section.

14.1.4 Create a Guest VM Server Object on Dell Compellent

Once all required fabric zoning changes have been made and saved to your fabrics, a guest VM server object can be created on the Dell Compellent Storage Center.

- 1) Launch **Dell Compellent Storage Center Manager**.
- 2) Create a new server folder (or navigate to an existing server folder) for your guest VM server object.
- 3) Right click on the server folder and select **Create Server**.

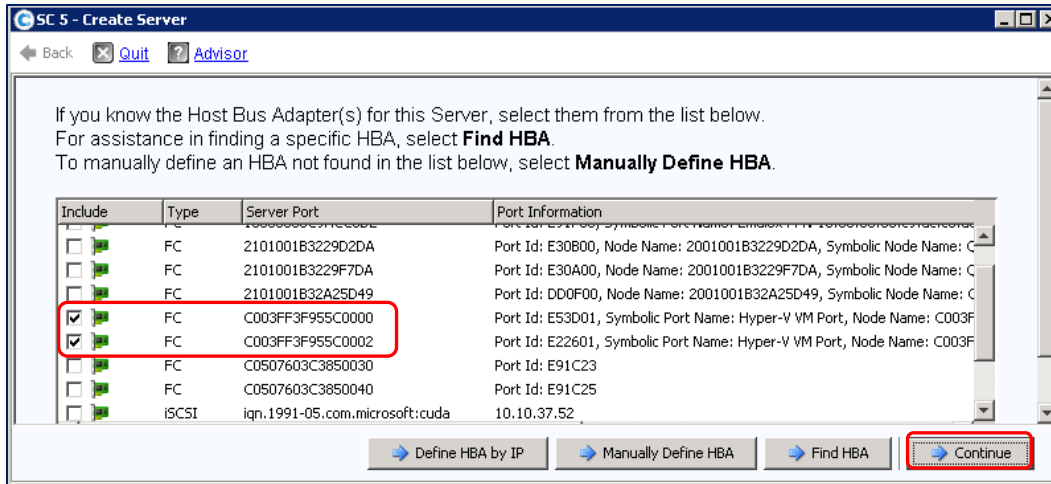


Figure 164. Select the Active WWPNS (Server Ports) for the Guest VM

- 4) Referring Figure 164 and Tables 7 and 8 above, the two currently active WWPNS (or Server Ports from the perspective of Storage Center) should be listed for this guest VM. The two inactive WWPNS (Server Ports) will not be listed (they will be mapped later).
- 5) Check the box for the two active WWPNS (Server Ports) as shown in Figure 164, and then click on **Continue**.

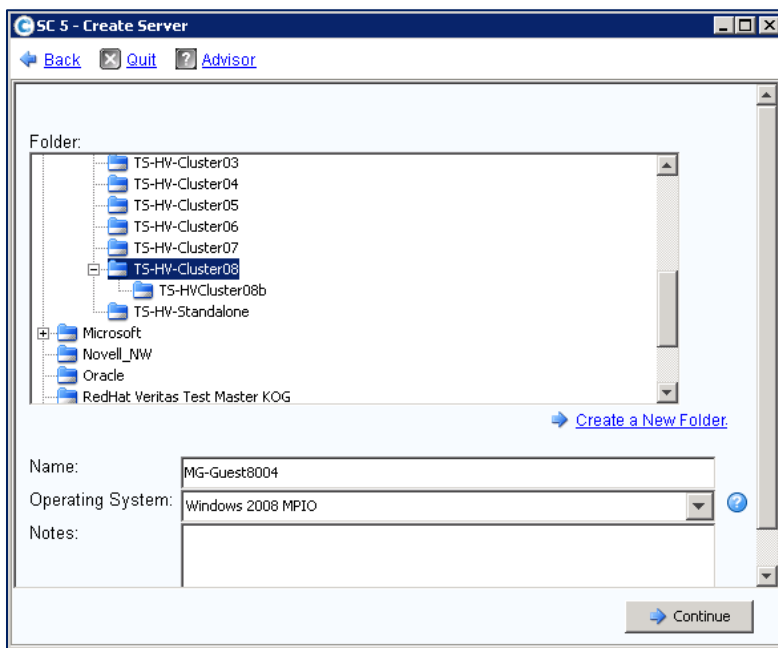


Figure 165. Provide a Name and OS for the New Guest VM

- 6) Provide a name for the server (**MG-Guest8004** in this example), select the desired operating system from the drop-down list (**Windows 2008 MPIO** in this example), then click on **Continue**.

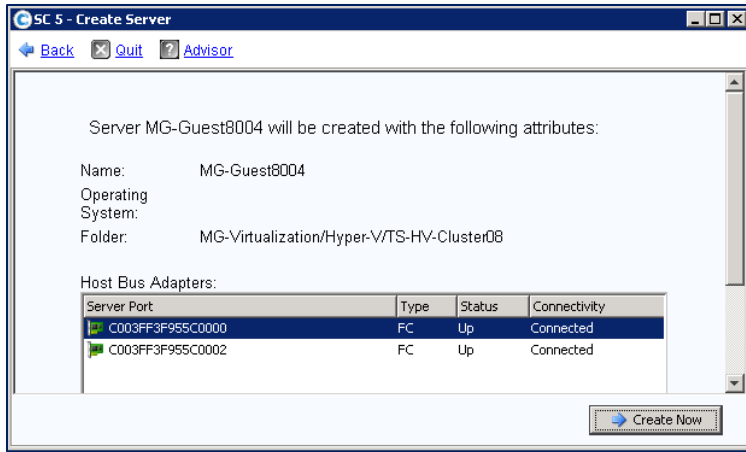


Figure 166. Verify Guest VM Settings Summary Screen

- 7) On the summary screen, verify that both ports are listed and then click on **Create Now**.
- 8) On the server creation confirmation screen, click on **Close** (storage will be mapped later).
- 9) Now that a guest VM server object has been created on Dell Compellent and the two active WWPNs (Server Ports) are mapped, the other two WWPNs (Server Ports) now need to be mapped. The best way to do this (to ensure reliable mappings, and to verify that the Cluster is configured correctly) is to fail the guest VM over to another node in the cluster, at which point the other set of WWPNs (Server Ports) will become active so they can be mapped.

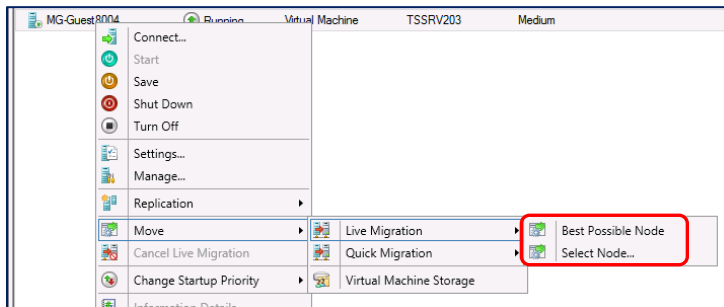


Figure 167. Live Migrate the Guest VM to Another Cluster Node

- 1) Using **Failover Cluster Manager** on a node in the cluster, live migrate the guest VM to another cluster node.
 - a. Right click on the guest VM.
 - b. Chose **Move**→**Live Migration**→**Best Possible Node** or **Select Node...**
 - c. Allow the guest VM to finish live migrating, and verify that it is now on another node in the cluster.

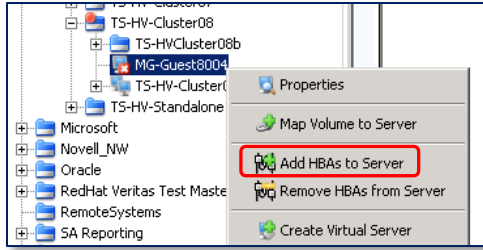


Figure 168. Add HBAs to the Guest VM

- Using **Storage Center Manager**, right click on the guest VM server object and select **Add HBAs to Server**.

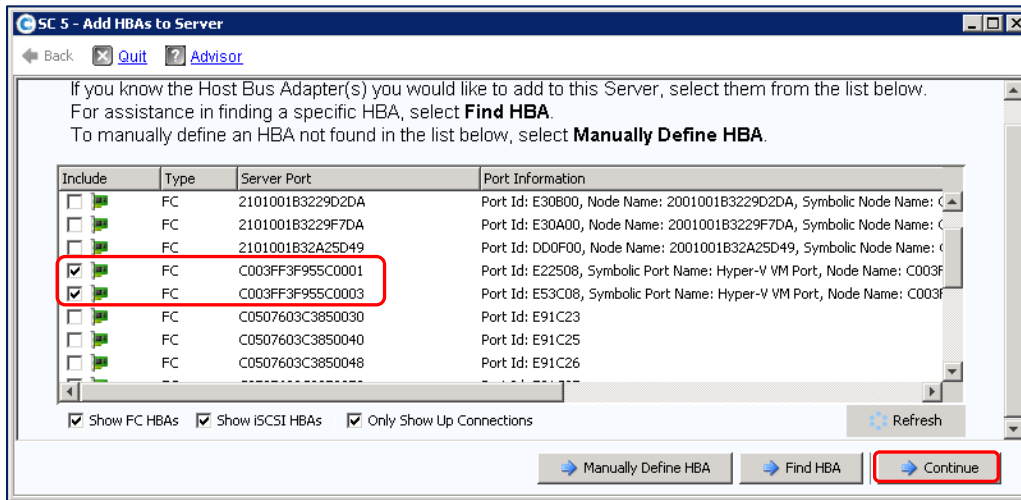


Figure 169. Map the Other Pair of WWPNs (Server Ports) to the Guest VM

- As shown in Figure 169 (and referring to Tables 7 and 8), the other pair of WWPNs (Server Ports) should now be listed since they are now active. Check the box in front of the two WWPNs (Server Ports) and then click on **Continue**, then on **Modify Now**.

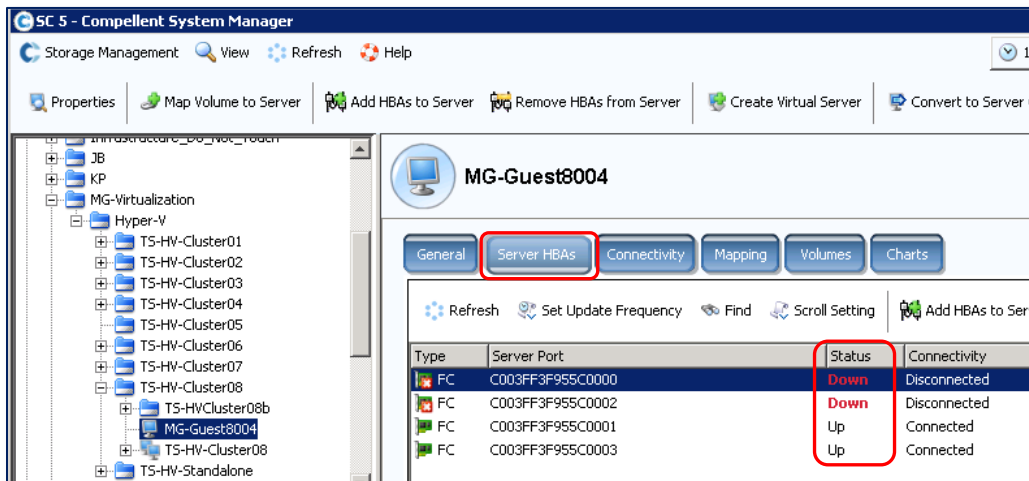


Figure 170. Guest VM Now Has All Four WWPNs (Server Ports) Mapped

- With the guest VM server object highlighted, under the **Server HBAs** tab, all four of the

guest VM's WWPNs (Server Ports) should now be listed: two that are up, and two that are down.

14.1.5 Create and Assign a SAN Volume to the Guest VM

The next step is to create and assign a SAN volume to the guest VM.

- 1) Create a new SAN volume (in the volume folder of your choice) on the Storage Center.
- 2) Select a storage profile, assign a Replay schedule, and configure Remote Replication for the volume as desired.
- 5) Right click on the new volume and select **Map Volume to a Server**.
- 6) Expand the **Servers** folder tree and click on the guest VM server object (**MG-Guest8004** in this example) and click on **Continue**.

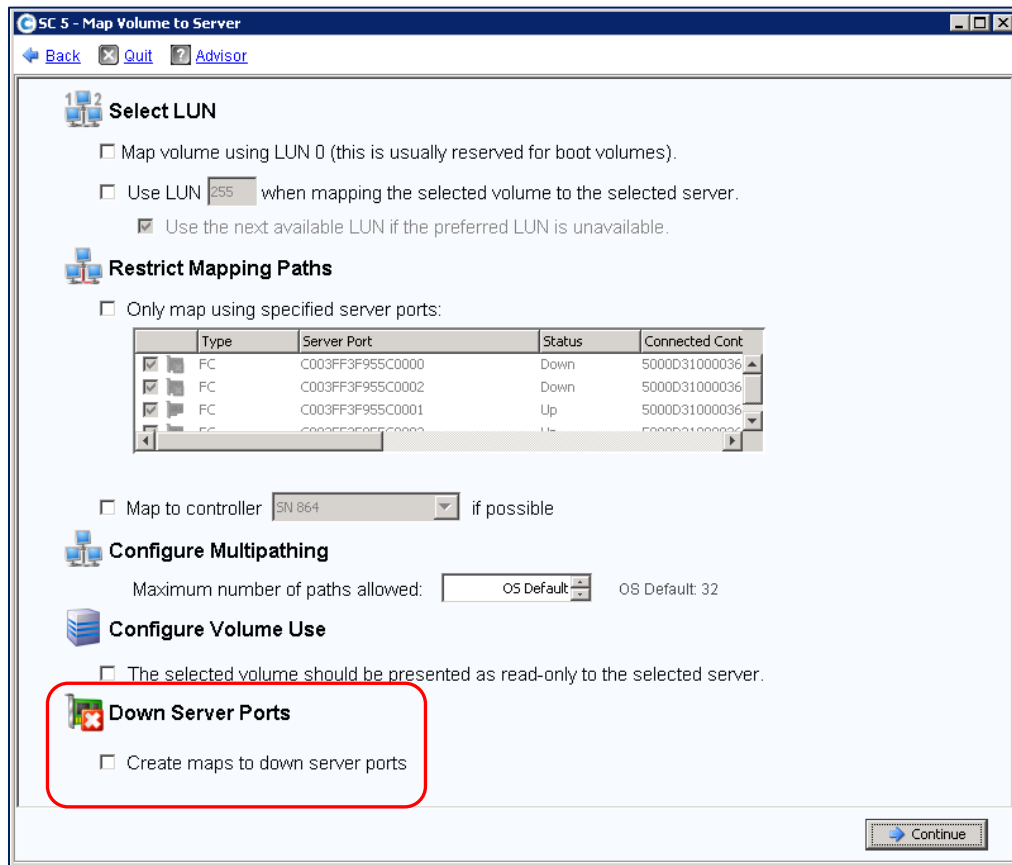


Figure 171. Map Storage to the Guest VM

- 7) Click on the **Advanced** button. In this example, no changes are made to the settings on this screen as the defaults are acceptable.
 - By default the Storage Center will map volumes to WWNs (Server Ports) that are "up" only. This is because in most circumstance, a server would never have "down" WWNs (Server Ports). However, with guest VMs that are using virtual fiber channel on a Server 2012 Hyper-V cluster, having "up" and "down" WWPNs (Server ports) is

unavoidable because it is by Microsoft design.

- While it is possible to force the wizard in Figure 171 to create mappings to the guest VM's "down" WWPNs (Server Ports) by checking the box **Create maps to down server ports**, leave this box unchecked. "Down" server HBA ports do not always map reliably or correctly, therefore, the "down" ports will be mapped later using a more reliable method.
- 8) Click on **Continue**, and then on **Create Now**.
 - 9) Now that the guest VM's two "up" WWPNs (Server ports) are mapped to the new SAN volume, use **Disk Manager** on the guest VM to perform a **Rescan Disks** from the **Action** menu and verify that the guest can now see the new SAN volume. If the new volume does not show up right away, wait a little while and run **Rescan Disk** again.
 - 10) The SAN volume may listed multiple times in Disk Manager. This will be corrected later when MPIO is configured.

Note: the new SAN volume must be visible to the guest VM (in Disk Manager) before continuing with the next steps.

- 11) The next step is to map the other two WWPNs (Server Ports) that are "down" to the SAN volume, so that it will have all four WWPNs (Server Ports) mapped.
- 12) Using **Failover Cluster Manager** on a node in the cluster, start a live migration operation on the guest VM to another cluster node. Because only two of the four WWPNs (Server Ports) are mapped to the SAN volume, the live migration operation will fail. However, during the time the live migration operation is attempting to complete, the other two WWPN's (Server Ports) will show as "up" allowing them to be mapped to the SAN volume.
 - a. Right click on the guest VM.
 - b. Chose **Move→Live Migration→Best Possible Node** or **Select Node...**
 - c. As soon as the live migration process starts, **Failover Cluster Manager** will bring the guest VM's other two WWPNs (Server Ports) online so they can be mapped to the SAN volume.

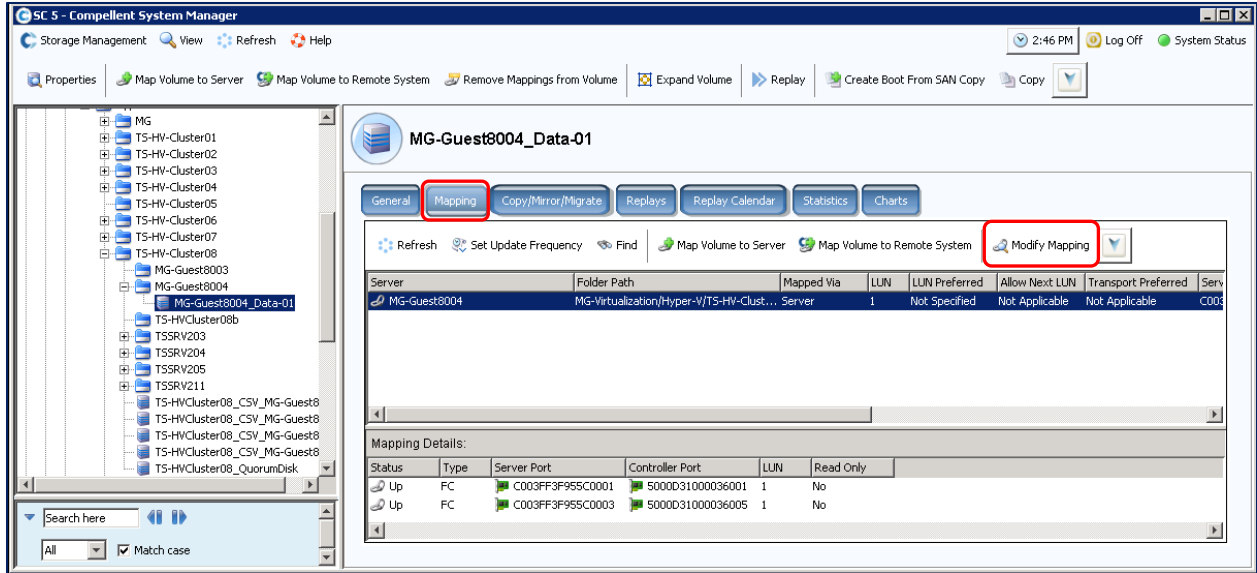


Figure 172. Modify the Guest VM's Data Volume Mappings

- d. Using **Storage Center Manager**, highlight the guest VM's data volume, and under the **Mapping** tab, click on **Modify Mapping**.
- e. Ignore the warning message and click on **Continue**.

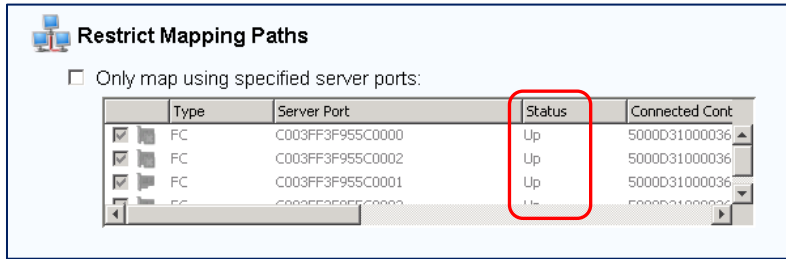
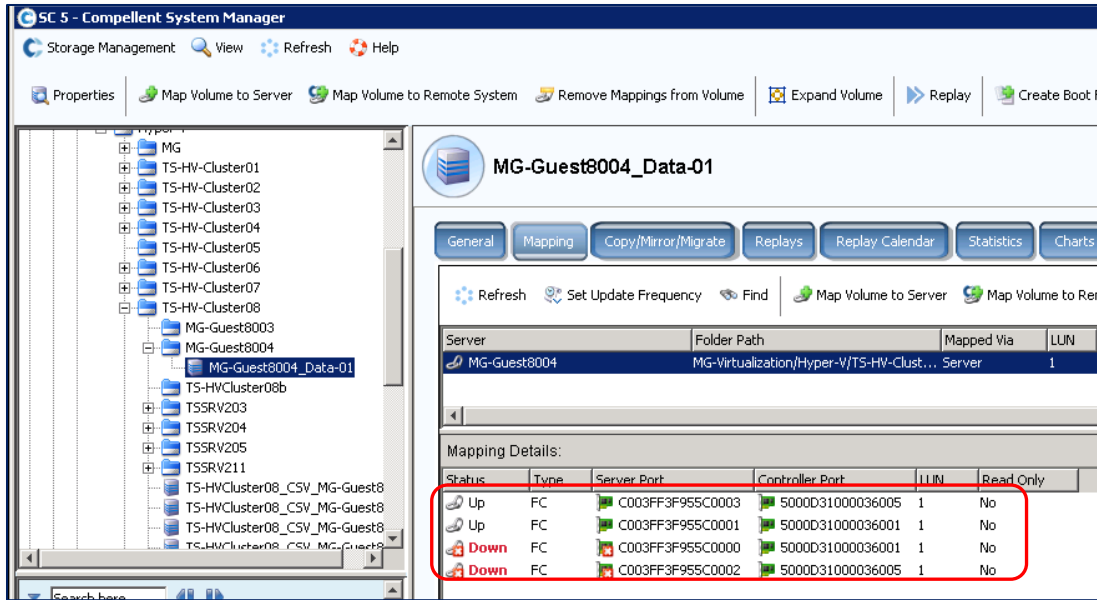


Figure 173. All 4 Guest VM WWPNS (Server Ports) Are Now Up Temporarily

- f. On the next screen, all four WWPNS (Server Ports) should now show as **Up** (temporarily while the live migration operation is running). Click on **Continue**, then on **Modify Now**. The additional two WWPNS (Server Ports) will be added to the SAN volume mapping because they are now up.
- g. Return to **Failover Cluster Manager** and right click on the guest VM, and select **Cancel Live Migration**. The live migration operation will eventually fail if it isn't canceled because the guest VM did not have all four WWPNS (Server Ports) mapped to the disk when the live migration operation was initiated.



- h. Storage Center Manager should now show four WWPNs (Server Ports) mapped to the guest VM's SAN volume: two that are up, and two that are down.
- i. On the guest VM, using **Disk Manager**, perform a **Rescan Disks** from the **Action** menu to refresh the disks and verify that it is still visible.
- j. Attempt to live-migrate the guest VM to another node in the cluster and it should complete successfully. If the guest does not live migrate, it may be necessary to reboot the guest VM to refresh it.

Status	Type	Server Port	Controller Port	LUN	Read Only
Down	FC	C003FF3F955C0003	5000D3100003600D	1	No
Down	FC	C003FF3F955C0001	5000D31000036009	1	No
Up	FC	C003FF3F955C0002	5000D31000036009	1	No
Up	FC	C003FF3F955C0000	5000D3100003600D	1	No

Figure 174. Guest VM Now Using the Other Two Server Ports

- k. After live migrating the guest VM, note that the other two WWPNs (Server Ports) will now show as active on the Storage Center.
- 13) The new volume may be listed multiple times on the guest VM if the Compellent DSM has not yet been installed. To install the DSM, open **MPIO Properties** on the guest VM and add the Compellent DSM, and then reboot the guest server if prompted. For more information on configuring MPIO, please see Section 4 of this document.
 - 14) After rebooting the guest VM, verify that the Compellent DSM is listed under the **MPIO Devices** tab of the **MPIO Properties** window, and that **Disk Manager** shows only one instance of the new volume on the guest VM. Configure MPIO settings as desired (set to round robin by default)
 - 15) Initialize the new volume, bring it on line, format it, and assign it a drive letter or mount point.

14.2 Guest VM Clustering with Dell Compellent

It is possible to configure services running on guest VMs (such as SQL) to be highly available by creating guest VM clusters (e.g. a virtual cluster on a physical cluster) when the guest VMs are mapped to direct-attached storage via iSCSI or virtual fiber channel. Table 9 below shows which options are currently supported by Dell Compellent.

Table 9. Virtual Clustering Comparison and Compatibility

Server OS	Virtual Guest Clustering Method	
	iSCSI	Virtual Fiber Channel
Server 2008 R2 Hyper-V	Supported	n/a
Server 2012 Hyper-V	Supported	Supported

- 1) To use cluster disks that are directly-attached to guests by virtual fiber channel to create a guest cluster, complete the steps in the previous section (Section 14.1) to configure 2 or more guests with cluster disks. Verify that the virtual fiber channel HBA WWNNs are unique for each guest VM, otherwise the Cluster Validation disk test for SCSI-3 Persistent Reservations will fail.
- 2) To use disks that are directly-attached to guests by using the Microsoft iSCSI initiator to create a guest cluster, complete the steps in Section 4.3 to configure two or more guests with cluster disks.

Note: Do not use Storage Center Manager to create cluster server objects for your clustered guests when using virtual fiber channel. Leave the server objects as individual objects. This is because it is not possible to include “down” HBAs when presenting a SAN volume to a server cluster object. This functionality will be included in a future release of the Storage Center Manager software. To work around this, simply assign cluster disks to each guest VM server manually, ensuring that when doing so, the cluster volume is assigned the same LUN number on each guest VM.

- 3) Once the guest VMs have cluster disks assigned, and all the required ports have been included in the mapping, use Failover Cluster Manager to verify that the guest VM servers pass cluster validation.
- 4) If the **iSCSI-3 Persistent Reservations** component of cluster validation fails, then ensure that the WWNN is unique for each guest VM’s virtual fiber channel HBA. By default, failover cluster manager assigns the same WWNN to all guest VM virtual fiber channel HBAs, but WWNNs must be different for each server when they are clustered.

Note: if it becomes necessary to assign a new WWNN to a guest VM, it is also necessary to change the WWPNs for the guest VM, and reconfigure any existing zoning for that guest VM. This is because the Storage Center caches WWNNs, and when a new WWNN is assigned to a

guest VM without also changing the WWPN, Storage Center will continue to use the cached WWNN until the Storage Center controllers are rebooted (which flushes the WWNN cache). A fix for this issue will be included with a future release of the Storage Center OS. To change the WWNN for a guest VM, please see Section 14.1.2, Step 9.

- 5) Assuming the cluster validation tests pass, create a server cluster and then proceed to configure the desired application or service (e.g. SQL) as highly-available.

15 Offline Data Transfer with Server 2012

Offline Data Transfer (ODX) is now supported (and enabled by default) with Server 2012, and will work out-of-the-box assuming that the underlying SAN storage also supports ODX. Dell Compellent support for ODX with Server 2012 will be introduced with the 6.3.1. release of the Dell Compellent Storage Center OS.

For more information, such as how enable or disable ODX, and how to establish performance benchmarks, please see this Microsoft reference:

<http://technet.microsoft.com/en-us/library/jj200627.aspx>

More information about ODX support with server 2012 will be included in the *Dell Compellent Best Practices Guide for Windows Server 2012*, due for release in Q1 of 2013.

16 Live Migration with Server 2012

With Server 2012, Hyper-V guests can now be migrated to not only another node in the same cluster, but to a different cluster, or a standalone node, or between standalone nodes.

Table 10 provides a quick comparison of the options and enhancements with live migration for Server 2008 R2 and Server 2012 Hyper-V. For more information about these features, please refer to the Microsoft document: *Feature Comparison – Windows Server 2008 R2 Hyper-V and Server 2012 Hyper-V* as listed in Section 19 – Additional Resources.

Table 10. Live Migration Features and Comparison

Hyper-V Version	Within the same cluster?	Between clusters?	Between standalone hosts?	Between standalone hosts and clusters?	Multiple guest VMs at the same time?	Guest Affinity? (synchronize two or more guest VM to live migrate at exactly the same time)
Server 2008 R2 Hyper-V	Yes	No	No	No	No	No
Server 2012 Hyper-V	Yes	Yes	Yes	Yes	Yes	Yes

The process of Live Migration between clusters and between standalone hosts is facilitated by the ability of Server 2012 Hyper-V to synchronize a guest's VHD/VHDX files, configuration, and memory (by using network bandwidth) to another host or cluster, while the guest is online. Once the data is all synchronized, the guest VM is cut over. While this new feature is SAN-agnostic, many of the best practices that have been reviewed above (e.g. thin provisioning, MPIO, use of Replays, etc.) still apply no matter where the guest VM's data resides.

Live migration of guest VMs between clusters or between standalone hosts may not always be realistic or practical due to time constraints. For guest VMs with small VHD/VHDX files, the live migration process is fairly quick. But, if the guest's VHD/VHDX files are extremely large (e.g. 50 TB), the time required to sync the data to another volume on a different host or cluster might be prohibitive.

When faced with moving a guest VM with extremely large VHD/VHDX files, it might be much quicker to use Data Instant Replay to present space-efficient View Volumes containing the guest VM's data to the desired new host or cluster. While this might involve a small amount of down time during a maintenance window (a few minutes) to move a large guest VM to another host or cluster, it might be a much more practical approach than waiting for hours or even days if moving 10's or 100's of TB of data.

In cases where there are multiple Storage Centers, Remote Replication of Volumes or LiveVolume can also be leveraged to move large data volumes to another storage center.

17 Hyper-V Replica with Server 2012

Server 2012 Hyper-V has made significant improvements with the disaster Recovery (DR) aspects of Hyper-V by introducing Hyper-V Replica. Hyper-V Replica allows administrators to replicate guest VMs to another location as part of a DR plan so that they can quickly be brought on line there. While Hyper-V Replica is storage-agnostic, many of the Dell Compellent features such as Replay Manager, Storage Center Replays, Live Volume, and Remote Replication will still continue to play an important and integral part of a comprehensive disaster avoidance and recovery plan for Microsoft Hyper-V.

More details about how Hyper-V Replica integrates with the Dell Compellent Storage Center and overall disaster avoidance and recovery strategy will be covered in the next release of the *Dell Compellent Disaster Recovery Best Practices Guide for Hyper-V*.

18 Conclusion

Hopefully this document has proved helpful and has accomplished its purpose by providing administrators with answers and best practices for many commonly asked questions associated with implementing Server 2008 and Server 2012 Hyper-V on Dell Compellent Storage.

With the release of Server 2012 Hyper-V, many new features were introduced. As additional best practices are developed for the new features, this guide will be updated accordingly.

19 Additional Resources

Below are some links to additional resources:

Microsoft Hyper-V (Server 2008) Planning and Deployment Guide:

<http://www.microsoft.com/downloads/details.aspx?familyid=5DA4058E-72CC-4B8D-BBB1-5E16A136EF42&displaylang=en>

Microsoft TechNet Hyper-V (Server 2008) document collection:

<http://technet.microsoft.com/en-us/library/cc753637.aspx>

Microsoft Technet Hyper-V (Server 2012) document collection:

<http://technet.microsoft.com/en-us/library/hh831531>

Feature Comparison – Windows Server 2008 R2 Hyper-V and Server 2012 Hyper-V

http://download.microsoft.com/download/2/C/A/2CA38362-37ED-4112-86A8-FDF14D5D4C9B/WS%202012%20Feature%20Comparison_Hyper-V.pdf

Microsoft Multipath I/O (MPIO) Users Guide for Windows Server 2012

<http://www.microsoft.com/en-us/download/details.aspx?id=30450>

Offloaded Data Transfers (ODX)

<http://technet.microsoft.com/en-us/library/jj200627.aspx>

Dell Compellent Documentation:

<http://knowledgecenter.compellent.com>