

EMC Symmetrix with Microsoft Hyper-V Virtualization

Applied Technology

Abstract

This white paper examines deployment and integration of a Microsoft Windows Server 2008 Hyper-V virtualization solution on EMC[®] Symmetrix[®] arrays. Details of integration with storage solutions as well as availability and mobility options for Windows Server 2008 Hyper-V are covered, including features provided in Windows Server 2008 R2.

October 2009

Copyright © 2009 EMC Corporation. All rights reserved.

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

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

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

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

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

Part Number h6643

Table of Contents

Executive summary	4
Introduction	4
Audience	5
Technology overview	5
Symmetrix overview	5
Microsoft Hyper-V hypervisor.....	6
Storage connectivity options for virtual machines	7
Child partition direct connectivity	7
Parent partition managed connectivity.....	8
Implementing Virtual Hard Disks	11
Implementing pass-thru disks.....	15
Storage connectivity summary.....	17
Availability and mobility for virtual machines	18
Windows Failover Clustering – parent partition	18
Windows Failover Clustering – child partition	20
Virtual machine Live Migration.....	21
Windows Failover Clustering Cluster Shared Volumes	22
Sizing of Cluster Shared Volumes	25
Site disaster protection with Windows Geographically Dispersed Clustering	25
Microsoft System Center Virtual Machine Manager	27
Integration with EMC layered technologies.....	31
EMC Solutions Enabler.....	31
Enhanced Virtual LUN Technology.....	33
Auto-provisioning Groups for Hyper-V.....	36
Auto-provisioning Groups and Microsoft Windows Failover Clusters.....	37
Connectivity recommendations.....	38
Conclusion	39

Executive summary

For many customers, there has been a growing need to provide ever-increasing physical server deployments to service business needs. This has subsequently led to a number of inefficiencies in operational areas. Servers have typically been overprovisioned in terms of CPU and memory resources in addition to storage resources. Power and cooling costs, as well as data center floor space requirements, grow with each additional physical server, whether the resources are overprovisioned or not. When considering the effects that large numbers of physical servers will have in an environment, and the inefficiencies of practices such as overprovisioning a physical server, the costs become exorbitant, and the resulting return on investment (ROI) is extremely poor.

Microsoft Windows Server 2008 Hyper-V provides customers with the ability to consolidate multiple physical server environments to achieve significant space, power, and cooling savings while maintaining availability and performance targets. EMC® Symmetrix® storage systems are able to provide additional value to customers by providing the ability to consolidate storage resources, implement advanced high-availability solutions, and provide seamless multi-site protection of customer data assets.

With the release of the Windows Server 2008 operating system, Microsoft introduced the latest in its series of server virtualization solutions. This hypervisor implementation is known as Hyper-V, and is available as a server role on the appropriately licensed Windows Server 2008 with Hyper-V SKU, and is also available with the Microsoft Windows Server 2008 R2 product. Additionally, the hypervisor functionality is available as a free download as Windows Hyper-V Server 2008 and Windows Hyper-V Server 2008 R2. The Hyper-V Server product is a bare metal install option. Hyper-V products are only available for the 64-bit (x64) release of Microsoft Windows Server 2008, and require that the server hardware platform supports hardware assisted virtualization (Intel-VT or AMD-V). As of Windows Server 2008 R2, Windows Server products are only available for x64 hardware. Support for the x86 platform ends with the Windows Server 2008 release.

As customers seek to consolidate data center operations, Microsoft's Hyper-V hypervisor provides a scalable solution for virtualization on the Windows Server platform. To further facilitate the cost savings available to customers that seek to consolidate operations, large-scale consolidation can benefit by optimizing and consolidating storage resources to a single storage repository. Additionally, many of the advanced features of the Hyper-V environment are either facilitated by, or enhanced with, the implementation of a scalable storage array.

Beyond providing protection and performance requirements through core system performance and RAID protection, Symmetrix arrays provide complementary technologies for Hyper-V environments that improve dynamic placement capabilities for Hyper-V landscapes. High availability and multi-site disaster protection can be transparently integrated to produce comprehensive solutions for customer deployments providing significant value-added solutions for consolidated Hyper-V deployments on Symmetrix storage systems.

EMC Symmetrix DMX™ and Symmetrix V-Max™ storage arrays are able to easily scale to the degrees required for large-scale consolidation efforts. With support of thousands of connected hosts, presentation of tens of thousands of logical units and advanced internal mechanisms such as snapshot and clone operations, and multi-site replication solutions to provide disaster restart/recovery solutions, Symmetrix systems are a central part of Windows Server consolidation efforts.

Introduction

This white paper presents an overview of Symmetrix and the Microsoft Hyper-V hypervisor. It provides storage connectivity choices for virtual machines, and explains availability and monitoring options for virtual machines. It concludes by describing how Symmetrix V-Max products like Solutions Enabler and Enhanced Virtual LUN Technology can integrate into virtual data center integrations.

Audience

This white paper is intended for Microsoft Windows Server 2008 administrators, storage architects, customers, and EMC field personnel who want to understand the implementation of Hyper-V solutions on EMC Symmetrix storage platforms.

Technology overview

Symmetrix overview

The EMC Symmetrix V-Max Series with Enginuity™ is the next generation of the Symmetrix product line. Built on the strategy of simple, intelligent, modular storage, it incorporates a new scalable fabric interconnect design that allows the storage array to seamlessly grow from an entry-level configuration into the world's largest storage system. Symmetrix V-Max arrays provide improved performance and scalability for demanding enterprise storage environments such as those found in large virtualization environments, while maintaining support for EMC's broad portfolio of platform software offerings.

The Enginuity operating environment for Symmetrix version 5874 is the latest Enginuity release supporting Symmetrix V-Max arrays. With the release of Enginuity 5874, Symmetrix V-Max systems now deliver new software capabilities that improve capacity utilization, ease of use, business continuity, and security. These features provide significant advantage to customer deployments in a virtualized environment, where ease of management, and protection of virtual machine assets and data assets are required.

Symmetrix V-Max arrays extend the scalability of previous generations of Symmetrix DMX technology, by providing a superior level of scalability, and support for a broad new range of drive technologies as detailed in Figure 1. Additionally, Symmetrix V-Max offers the ultimate in flexibility, including the ability to incrementally increase back-end performance by adding V-Max Engines and storage bays. Each high-availability V-Max Engine controls eight redundant Fibre Channel loops that support up to either 240 or 360 drives depending upon configuration. Subsequently, each high-availability V-Max Engine provides front-end as well as back-end connectivity, providing enhanced scalability.

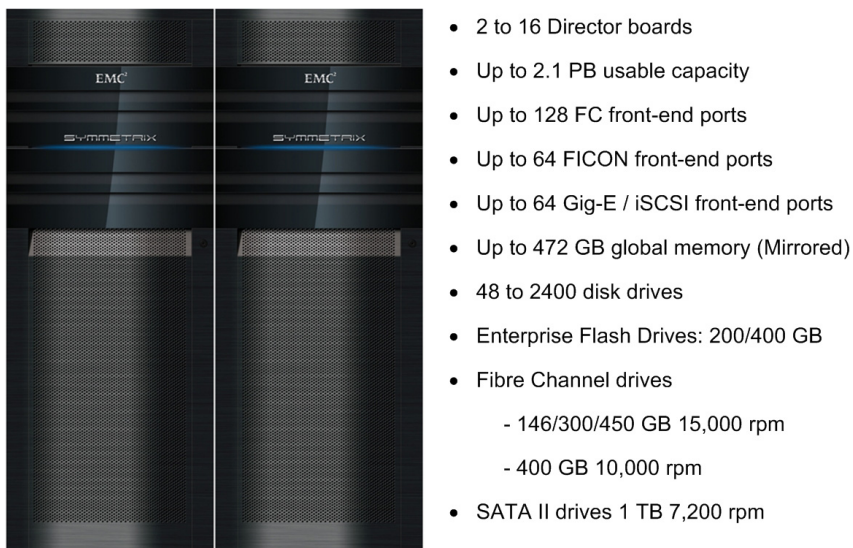


Figure 1. Symmetrix V-Max hardware scalability

The Symmetrix V-Max systems also maintain customer expectations for high-end storage in terms of availability. High-end availability is more than just redundancy; it means nondisruptive operations and upgrades, and being “always online.” Beyond previous Symmetrix generations, Symmetrix V-Max arrays provide:

- Nondisruptive expansion of capacity and performance at a lower price point

- Sophisticated migration for multiple storage tiers within the array
- The power to maintain service levels and functionality as consolidation grows
- Simplified control for provisioning in complex environments

Microsoft Hyper-V hypervisor

Microsoft Windows Server 2008 provides the Hyper-V server role on the applicable versions of Windows Server. In the initial release of Windows Server 2008, separate product releases included the Hyper-V role, and were required to be ordered explicitly. When a Windows Server instance has the Hyper-V role installed, the original operating system instance is referred to as the “parent partition.”

When the Hyper-V server role is installed, the Windows Hyper-V virtualization hypervisor is installed for the parent partition. Utilizing the functionality implemented by the hypervisor, and managed through the Hyper-V Manager Management Console (MMC) shown in Figure 2, it is possible to define virtual machine instances.

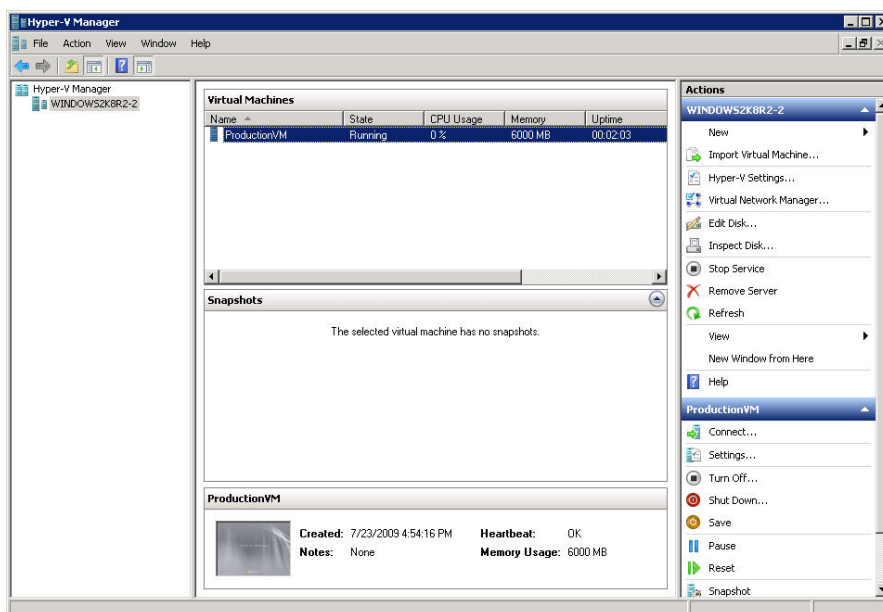


Figure 2. Hyper-V Manager Management Console

In more complicated Hyper-V deployments that may be comprised of a large number of physical servers, and a multitude of virtual machine instances, customers may choose to implement products like Microsoft System Center Virtual Machine Manager 2008 (SCVMM). The SCVMM solution provides a comprehensive management framework with centralized command and control features. SCVMM also includes additional functionality in the form of the Performance and Resource Optimization (PRO) subsystem. The PRO functionality has a dependency on Microsoft System Center Operations Manager (SCOM), allowing customers to build automatic and dynamic management capabilities into a Hyper-V landscape. Such configurations may allow for dynamic placement of virtual machine resources based on changing characteristics of the data center. More information regarding SCVMM and its integrated functionality can be found at <http://www.microsoft.com/systemcenter>.

A virtual machine instance is typically identified by a configuration file defining the configuration of the virtual machine. The virtual machine instance in Hyper-V environments is also often referred to as a “child partition.” The definition of the child partition includes such aspects as processor count, memory configuration, network connectivity, and other hardware details, as well as a complement of one or more storage devices that represent the storage resources utilized by the operating system instance. These features are configurable through the child partition settings options in the Hyper-V MMC, as shown in Figure 3.

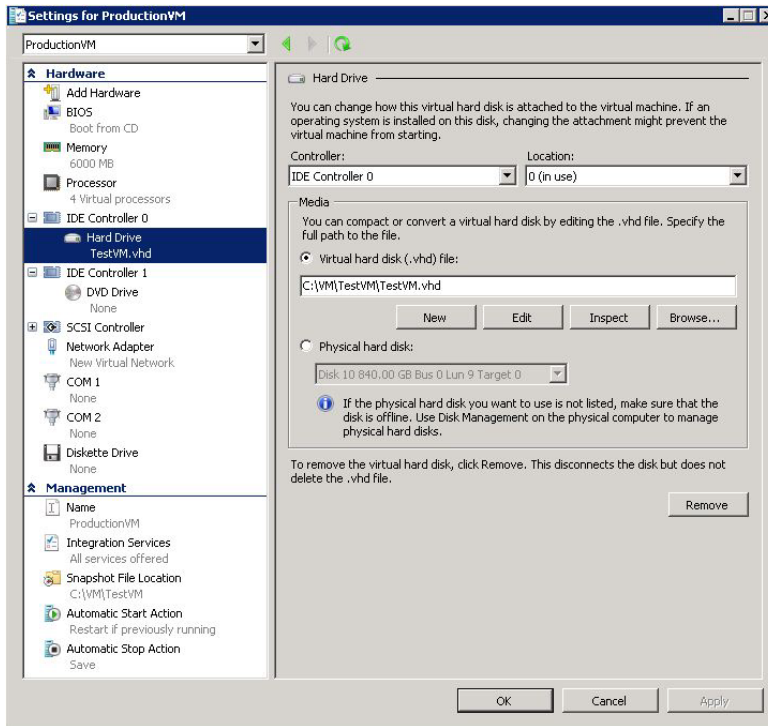


Figure 3. Virtual machine/child partition configuration settings

Two primary representations of storage devices may be configured for a virtual machine from the settings options. Either a storage device will be provisioned as a Virtual Hard Disk (VHD) and connected to one of the IDE or SCSI Controller adapters, or the device will be connected as a physical hard disk (a pass-thru storage device).

Storage connectivity options for virtual machines

Microsoft Hyper-V configurations support multiple deployment models for connectivity to Symmetrix arrays. However, there are two basic methods of a virtual machine itself gaining access to storage resources. These two forms may be characterized as either connectivity directly to the virtual machine (via a network resource) or connectivity provisioned via the parent partition. In this latter form of connectivity, the storage device is managed by the parent partition but utilized by the virtual machines in some manner. In the case of direct connectivity from the virtual machine, the virtual machine accesses storage via an iSCSI connection utilizing network connectivity and the parent partition is not physically involved in managing the storage device. It is also possible, via network connectivity, to access storage resources via SMB, but this style of connectivity is not covered in this white paper. It should be noted, however, that although iSCSI connectivity may be described as “direct” technically, the network connectivity is occurring over virtualized network adapters managed by the parent.

Microsoft provides, supports, and recommends running a Hyper-V parent in the minimal footprint of a Windows Server 2008 Core installation. Additionally, the usage of the Hyper-V Server 2008 bare metal installation option is Microsoft supported and recommended. This white paper only describes the usage of the Windows Server 2008 Full installation option.

Child partition direct connectivity

Virtual machine instances running Windows Server 2003 or Windows Server 2008 and Windows Server 2008 R2 are able to utilize storage provided directly to the virtual machine as an iSCSI target. In this form of connectivity, the hosted operating system of the virtual machine will need to implement the Microsoft iSCSI Initiator software and be able to access network resources through a virtual network interface.

As the virtual machine itself is directly accessing the iSCSI storage device via the network, the operating system within the virtual machine is responsible for all management of the disk device and subsequent volume management. An iSCSI target device will need to be appropriately configured for the virtual machine to access the iSCSI devices. For further information on configuration of the Microsoft iSCSI Initiator, refer to the *EMC Host Connectivity Guide for Windows* available on Powerlink®.

For most configurations it will still be necessary to initially provision a VHD device to support the installation of the virtual machine operating system. In this case, the creation of that VHD device will proceed in the manner outlined in the “Parent partition managed connectivity” section.

Third-party hardware iSCSI solutions may support a Boot from iSCSI SAN solution; however, these solutions are beyond the scope of this white paper due to the specific details required for each implementation.

Parent partition managed connectivity

When initially deploying a child partition under Hyper-V, it will commonly be necessary to provide the location for the VHD storage that will represent the operating system image. For the majority of virtual machine configurations, this will be the classic installation method. When defining a new virtual machine, the initial set will require that the virtual machine is provided a logical name such that it can be identified within the environment. This name does not control the operating system name that will be defined when the virtual machine is configured, but is displayed within the Hyper-V management products. As shown in Figure 4 the initial configuration requires a Hyper-V management name, and also the location for the virtual machine configuration information, which is highlighted. If the intent is to provide any form of high availability for the virtual machine, then this location should represent a SAN device that will be available to all high-availability nodes.

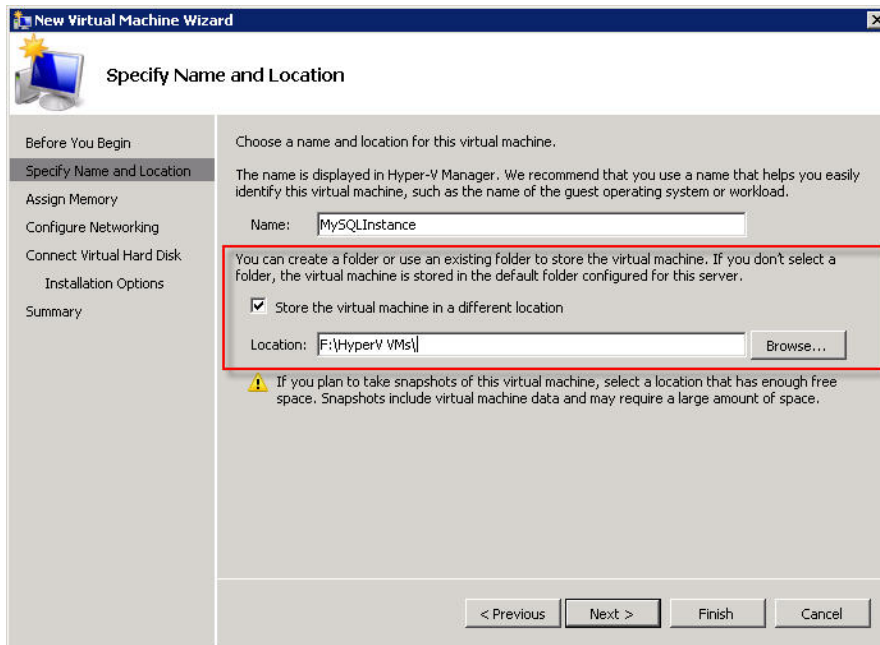


Figure 4. Create a virtual machine

Subsequent steps in the configuration wizard will request sizing information for memory allocation and network connectivity, which are beyond the scope of this white paper. For guidance on these parameters please refer to the Hyper-V online documentation.

The New Virtual Machine Wizard will proceed to the configuration of the Virtual Hard Disk (VHD) for the operating system installation as shown in Figure 5. The default location for the VHD will be based on the previous location specified in the Location field, and the VHD Name field will be based on the name provided for the virtual machine.

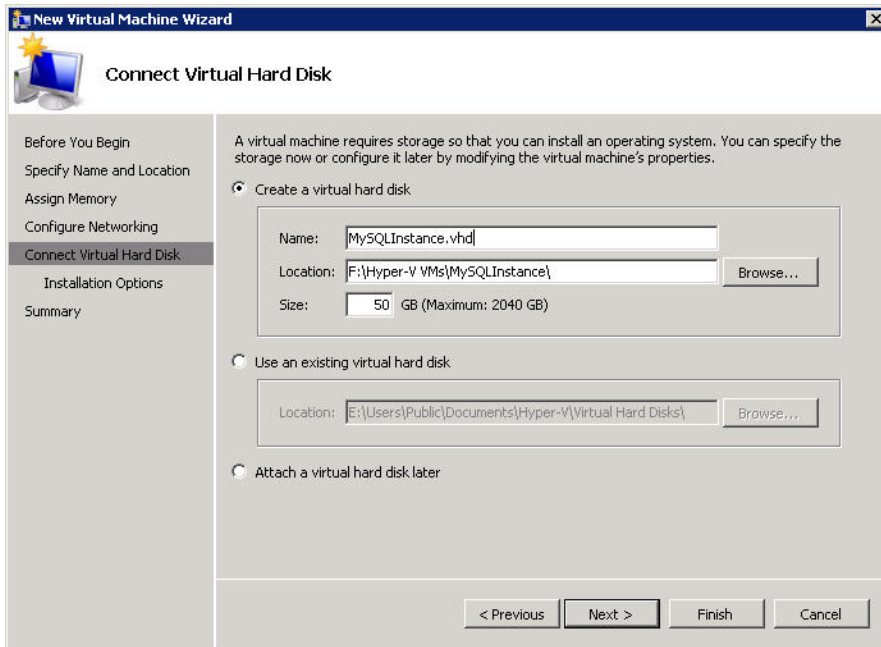


Figure 5. Default definition of the Virtual Hard Drive

The VHD should be sized appropriately for the operating system being installed. When configured through the wizard in this manner, the VHD created will be a Dynamic VHD file. It is possible to manually configure VHD devices for the virtual machine so as to specify specific VHD characteristics. To allow for manual configuration of VHD devices, if required, select the **Attach a virtual hard disk later** radio button. Subsequent sections will detail the options available for the manual option, and thus will assume that the option to attach a VHD at a later time was taken.

Once the New Virtual Machine Wizard has completed, selecting the settings option from the Hyper-V Manager console will allow for further modifications to the virtual machine configuration. The configuration is stored in an XML document located in the virtual machine directory, as shown in Figure 6. The name of the configuration file will be based on a Global Unique Identifier (GUID) for the virtual machine.

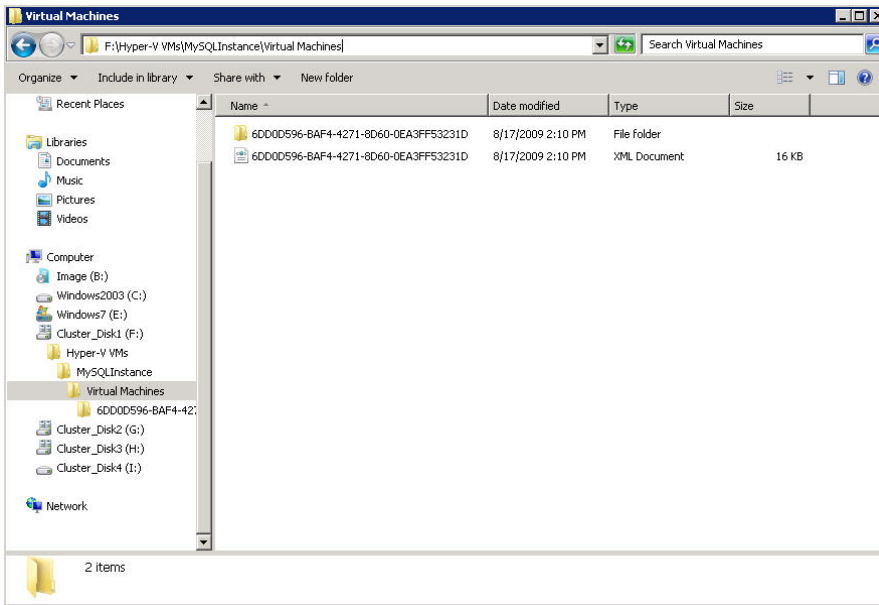


Figure 6. Virtual machine configuration file

In Figure 7 details for the newly created virtual machine can be seen. Note that currently no VHD devices have been configured, and so no such devices appear in the list. By default, a VHD defined by the New Virtual Machine Wizard would create a VHD that would be mapped to “IDE Controller 0”.

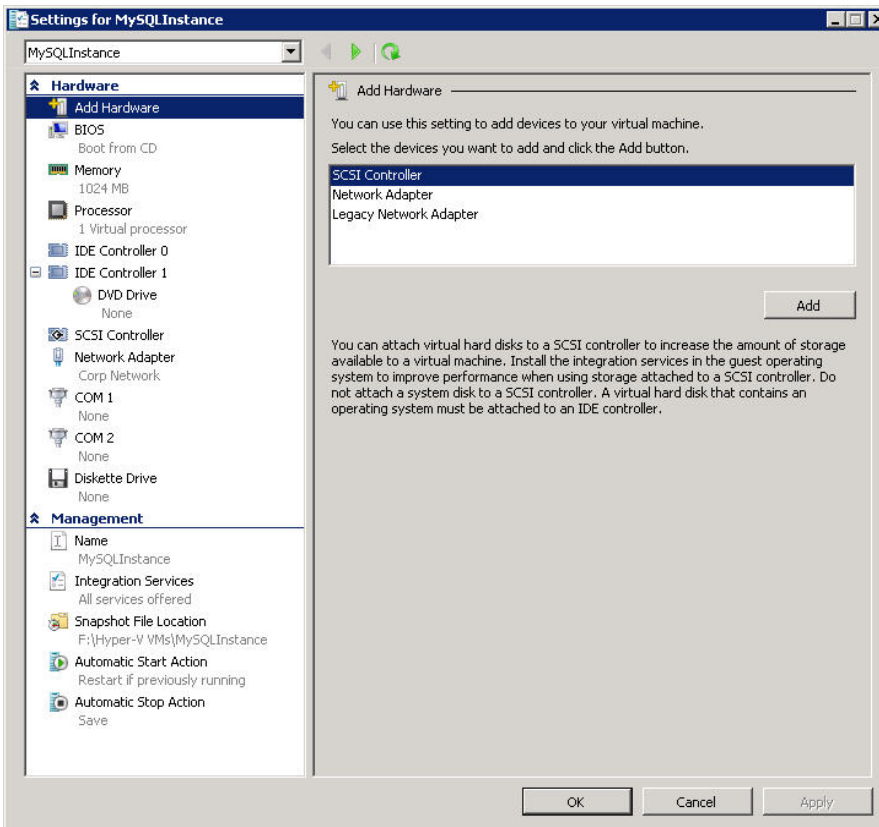


Figure 7. Virtual Machine configuration settings

Manually configured VHDs may be mapped to either the IDE or SCSI controllers defined in the virtual machine configuration. At least one such VHD must typically exist to host and install the operating system.

Implementing Virtual Hard Disks

It is possible to define VHD devices that may be later repurposed to virtual machines. Creation of the VHD devices may be done outside of the context of a virtual machine, by selecting the New Hard Disk options available within Hyper-V Manager, as shown in Figure 8. The overlaid arrows point to the two areas where the new VHD options may be selected.

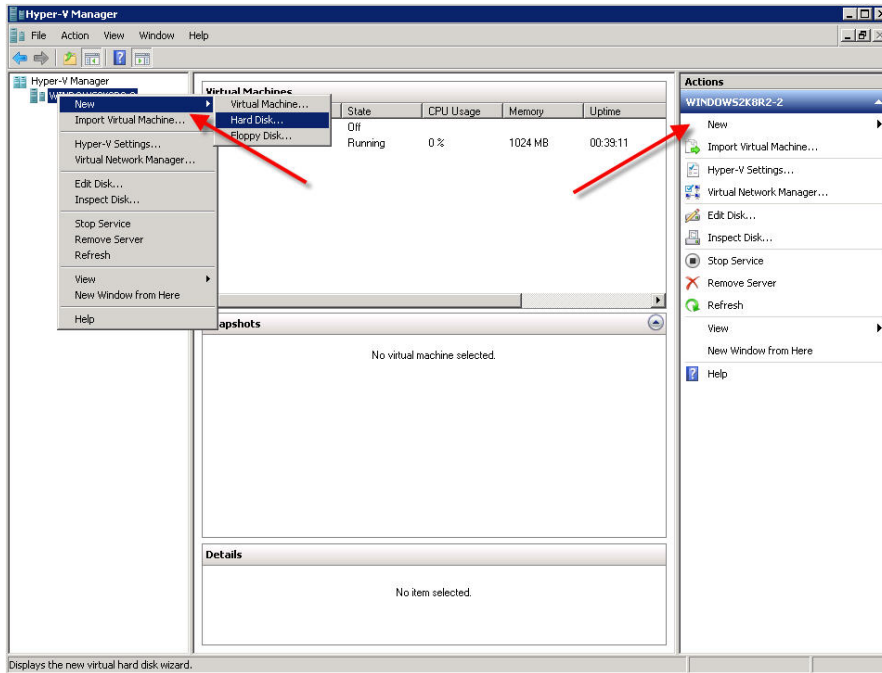


Figure 8. Defining new VHD devices

It is also possible to define and associated a VHD from within the settings option for a virtual machine by selecting the controller (IDE or SCSI) to which a VHD will be associated. To define and map the VHD within the virtual machine, open the settings for the virtual machine as shown in Figure 9. By selecting the hardware device, in this example “IDE Controller 0” it is possible to define the new VHD to be created and assigned to that controller.

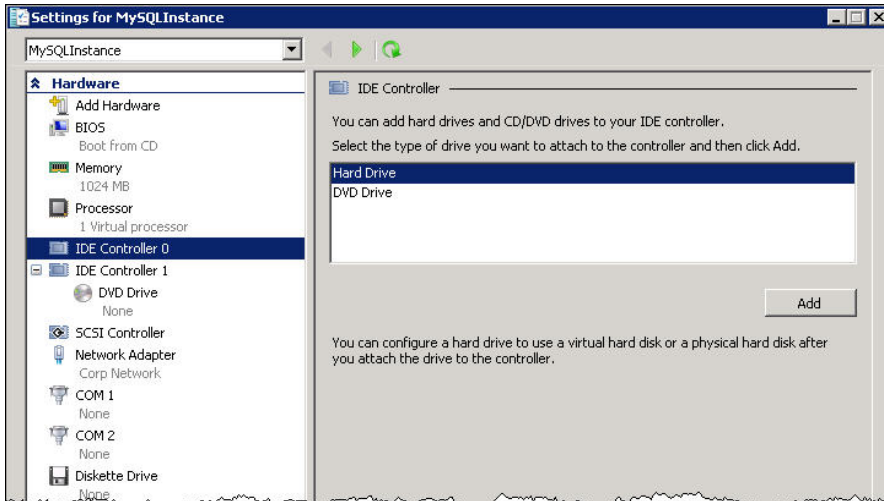


Figure 9. Create a new VHD within virtual machine settings

After selecting **Hard Drive** as the device to be created, and clicking the **Add** button, it is possible to define the various options for the new VHD to be created. Careful consideration should be given to the location for the new VHD device, specifically in those cases where the virtual machine is to be provided high availability through clustering functionality, the VHD must be located on SAN devices.

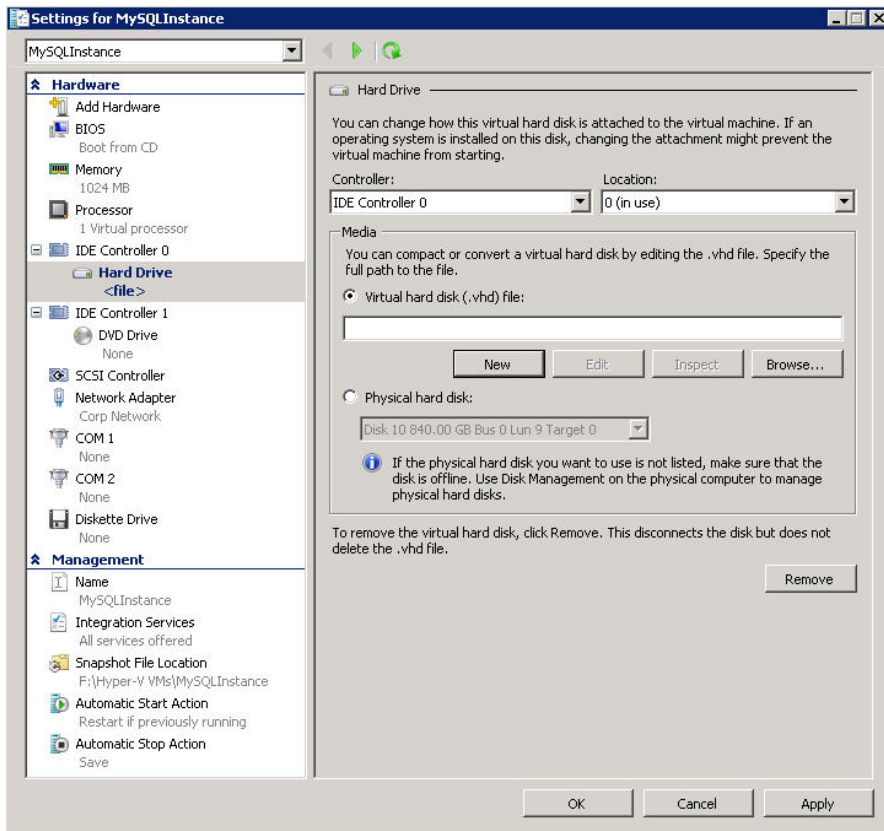


Figure 10. Create a new VHD within virtual machine settings

When defining the new VHD and assigning it to the virtual machine after selecting the appropriate controller, the definition must specify those required settings for the controller in use. For example, when defining a VHD for an IDE controller, as shown in Figure 10, it is necessary to specify the primary IDE

controller to be used (only two IDE controllers are supported), and the IDE device on the IDE controller. Only two devices per IDE controller are supported, which subsequently means that only four IDE VHD devices can be configured for any given virtual machine. If additional VHD devices are required, they must be defined as SCSI controller managed devices.

SCSI controllers are able to support multiple disk devices per controller, and are therefore a much more scalable solution for configurations when multiple LUN devices or multiple VHD devices are required. Such configurations may be required for very large storage requirements for a virtual machine. Each VHD has a maximum size of 2,048 GB irrespective of which controller it is located on.

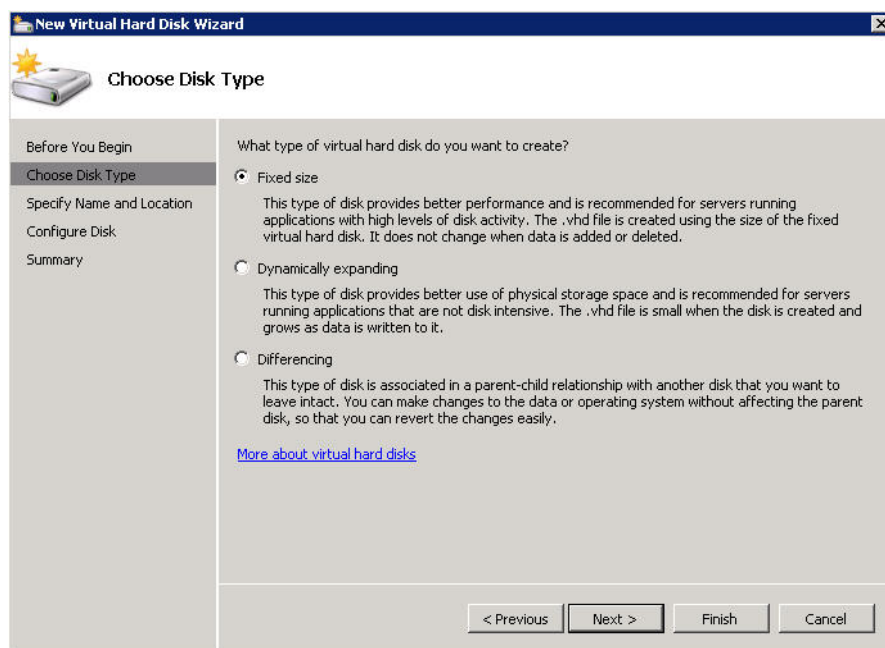


Figure 11. VHD device types

Three different types of VHD disks are available when configuring new or additional storage devices, as shown in Figure 11. The choice between a “Fixed size” and “Dynamically expanding” format will generally be based on the storage utilization requirements, in as far as there is a difference in how storage is allocated for these two types. For this reason, the two selections will subsequently affect storage provisioning functionality such as that provided by Symmetrix Virtual Provisioning™. A Fixed sized VHD device will be fully written to at creation time, and as a result, when selecting this VHD type, the creation of the device can take a considerable amount of time.

Dynamically expanding VHD devices will not preallocate all storage defined for them; however these devices may suffer a slight degradation in performance as a result of the need to allocate actual storage when the operating system or applications within the virtual machine require more allocations. These storage allocations are simply those to acquire additional storage from the parent partition for the VHD device – the usage of Virtual Provisioning does not affect this behavior.

Third-party tools exist in the public domain that allow for the creation of a fixed VHD device that behaves similar to a Dynamic VHD in as far as preallocation is not executed. Such solutions are functional, but may not be entirely supported by Microsoft. One such tool is provided at <http://code.msdn.microsoft.com/vhdttool>. Customers should consider the support implications with such public domain solutions.

When locating the VHD, it is possible to define any storage location accessible from the parent partition. In this way it is possible to define storage locations on separate LUNs visible to the parent partition, and thereby distribute the workload generated by the virtual machine and its applications. The definition of the location is specified by the subsequent dialog as shown in Figure 12. If the desire is to create a configuration that provides a high-availability solution for the virtual machine, then the VHD should be located on a SAN device.

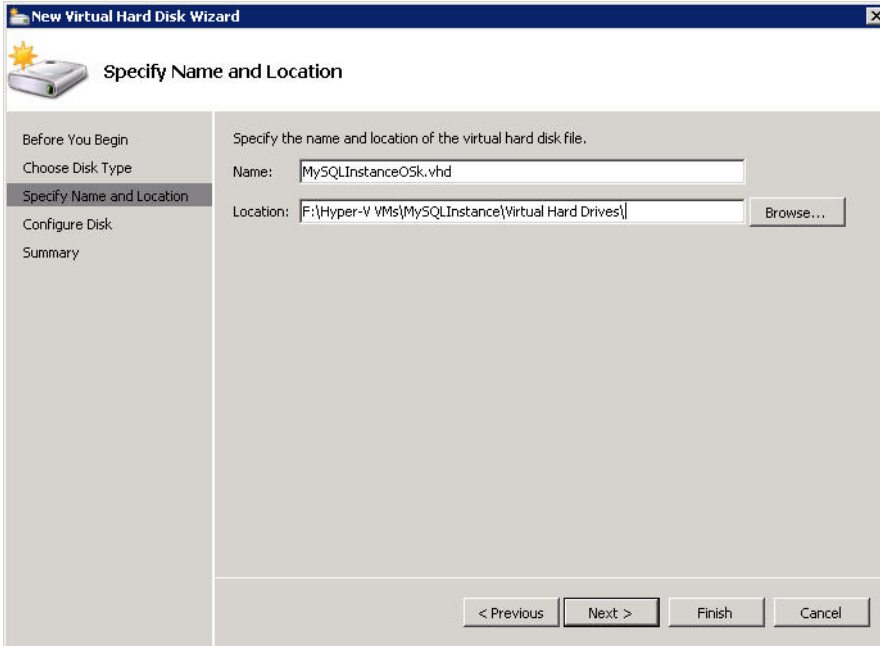


Figure 12. Definition of the location of the VHD

The final step in the creation of the VHD device is to specify the size, as shown in Figure 13. The maximum size supported by a single VHD device is 2,048 GB. As discussed, when more storage space than a single VHD device is required for a virtual machine, multiple VHD devices may be presented to the virtual machine, either through the remaining IDE device locations, or via one of the SCSI controllers.

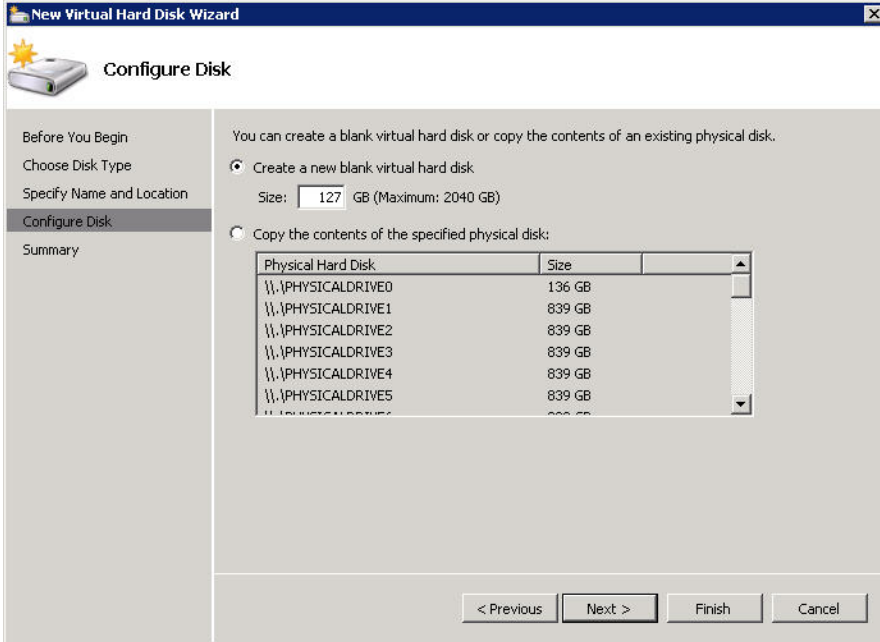


Figure 13. Specification of the VHD size and copy option

Once complete, the New Virtual Hard Disk Wizard will create and associate the new VHD device to the IDE or SCSI port that was specified. Figure 14 demonstrates the configuration settings after the creation of the fixed VHD device, and its assignment to IDE Controller 0 at Location 0.

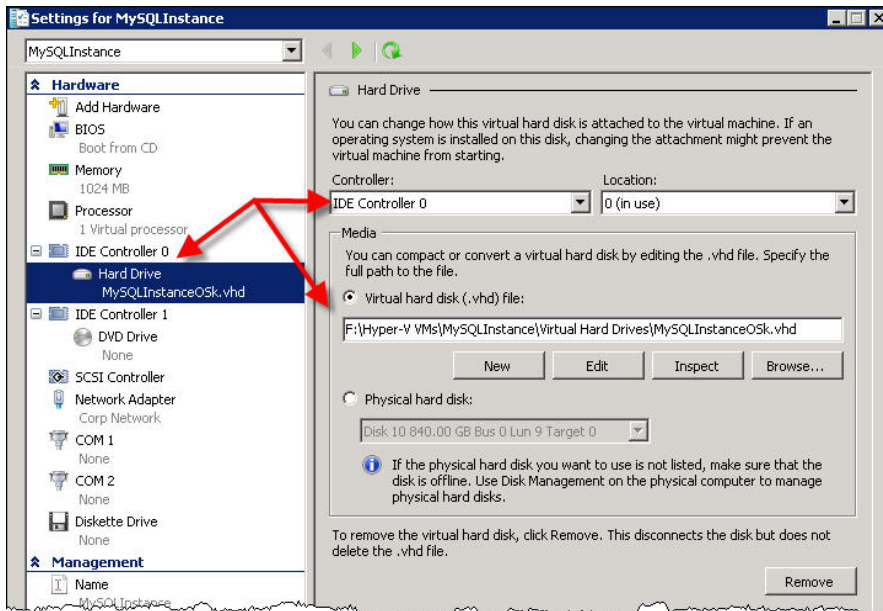


Figure 14. Settings after VHD creation

The third VHD format option available is that of a Differencing Disk. As detailed in the dialog box in Figure 11, this disk device is configured to be used to provide an associated storage area created against a source VHD. This style of configuration may be used when, for example, a Gold Master VHD is created, and is being used to create multiple virtual machine instances. In such configurations, it will be necessary to protect the Gold Master from being updated by any individual virtual machine. But each virtual machine will necessarily have to write its own changes. In this instance, the Gold Master VHD would behave as a read-only device, and all changes written by the virtual machine would be saved to the Differencing Disk device. There will always be an association between the Differencing Disk and the Gold Master, for without the original Gold Master, the Differencing Disk only maintains changes, and does not represent a fully independent copy.

Implementing pass-thru disks

Due to the nature of the way that I/O generated to a VHD device located on a volume managed by the parent partition is processed, a number of levels of indirection are imposed. I/O within the virtual machine is serviced by the virtual machine's operating system, and passed to the storage device. In turn, as the VHD is physically owned by the parent partition, the parent must now receive and re-drive the I/O to the physical disk that it owns. This multi-level indirection of I/O does not provide the best performance, although the overhead is relatively small. Invariably, the best performance will be obtained when there are the least levels of indirection. For storage devices presented to a virtual machine, this will mean that the best performance may be obtained by utilizing pass-thru devices or iSCSI devices directly to the virtual machine.

Pass-thru devices are required to be configured as offline to the parent partition and will therefore be inaccessible for any parent managed functions, such as creation or management of volumes. These offline disk devices are then configured as storage devices directly to the child virtual machine using the New Virtual Hard Disk Wizard. The disk management console or alternatively, the DISKPART command line interface, can be used to transition SAN storage devices between online and offline status. Figure 15 shows disk status displayed through the disk management console. To switch a disk between online and offline state, the context menu (right-click) when selecting the disk device can be used. By default, SAN devices will be presented in an offline state. The default storage devices state is managed by the SAN Policy set via the DISKPART command line interface.

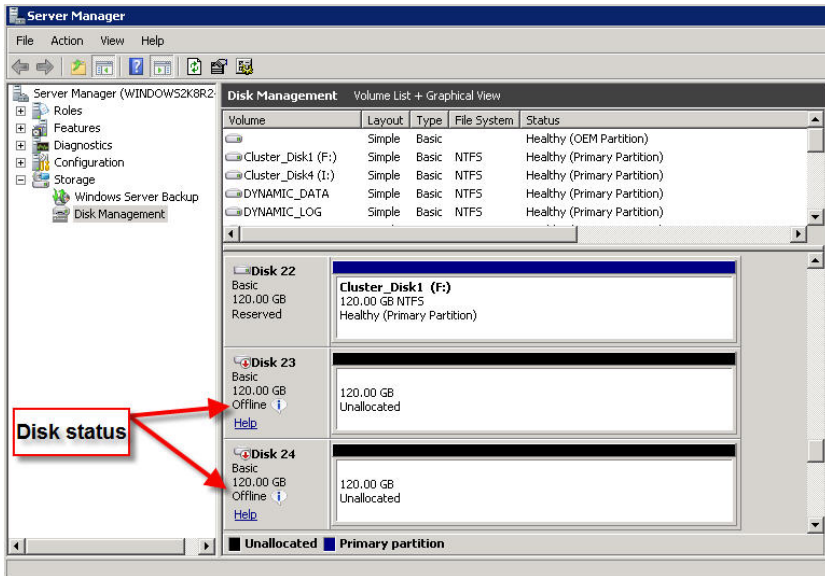


Figure 15. Disk status shown by Disk Management

Utilizing the Hyper-V MMC, it is possible to configure pass-thru devices as shown in Figure 16. Disks may be allocated against a SCSI controller configured for the virtual machine. Each SCSI controller is able to map up to 64 pass-thru devices, and up to four discrete SCSI controllers may be configured to an individual virtual machine. This provides support for up to 256 SCSI devices.

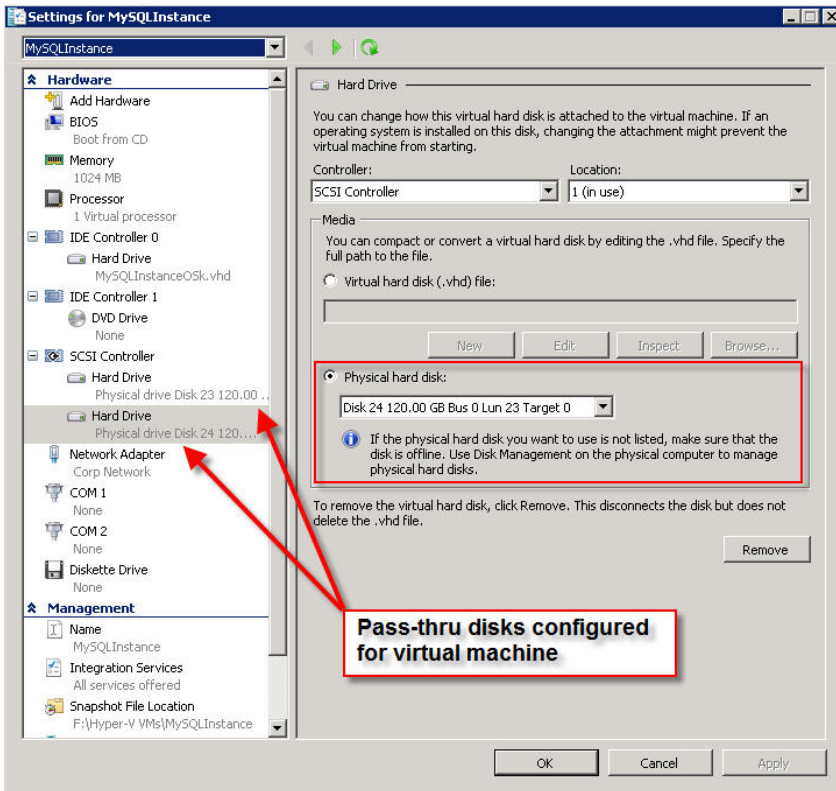


Figure 16. Configuration of pass-thru disk devices for a virtual machine

Once the required disk devices have been configured as pass-thru devices to the virtual machine, the operating system of the virtual machine will detect and display them as shown in Figure 17. In this

instance, the virtual machine has been configured with a VHD device that is used as a boot device. This boot device is displayed as the “Virtual HD ATA Device”. The two pass-thru devices configured are shown as “EMC SYMMETRIX SCSI Disk Device” as this is the detected storage device from the Windows Server 2008 R2 operating system of the virtual machine.

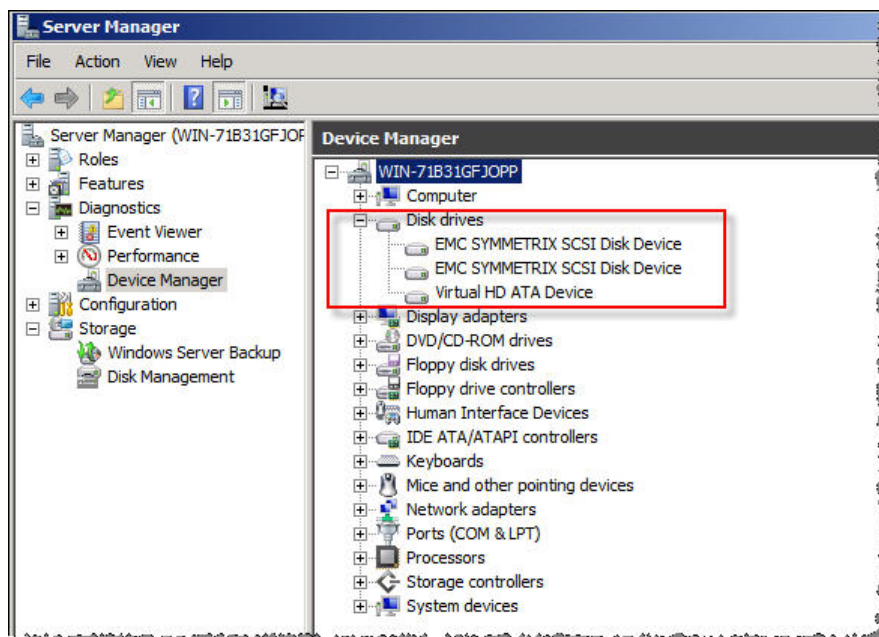


Figure 17. Disk devices within a virtual machine

Storage devices configured as pass-thru devices to a virtual machine will need to be configured in the same manner as is typical of storage devices to a physical server. Administrators should follow the recommendations provided for a physical environment, which may include the requirements to align partitions on applicable operating systems. In the case of Windows Server 2008 and Windows Server 2008 R2, manual partition alignment is not required, as partitions are automatically aligned to a 1 MB offset. The creation of the NTFS volume should also follow recommendations, such as selecting an Allocation Unit size of 64 KB when formatting volumes to be used for applications such as Microsoft SQL Server and Microsoft Exchange Server.

It is possible to deploy virtual machine instances that utilize pass-thru devices as the boot device for the operating system disk device. In such cases, it will be necessary to define the pass-thru device prior to installing the operating system of the virtual machine, and selecting the pass-thru disk (configured via the IDE controller) as the install location.

Storage connectivity summary

Symmetrix storage arrays provide and support all forms of storage connectivity required by Windows Server 2008 Hyper-V. Deployments may select any form of parent managed, or direct to child storage connectivity, and solutions may even combine multiple forms of connectivity to satisfy application-level requirements.

Each form of storage connectivity provides differing management or operational features. For example, storage that is provisioned directly to a child partition either using iSCSI connectivity or using pass-thru disks restricts that storage volume to being utilized by the virtual machine exclusively. Conversely, storage allocated in the form of VHD devices created on volumes within the parent partition will allow for a single LUN to be shared among any number of virtual machines, by allowing for the various VHD devices to be co-located on the parent-managed volume.

Utilizing a common parent-managed volume to co-locate VHD devices can also affect some high-availability or mobility solutions, since a change to the single LUN will affect all virtual machines located on the LUN. This may affect configurations using Windows Server 2008 Failover Clustering. However, the implementation of Clustered Shared Volumes (CSV) with Windows Server 2008 R2 Failover Clustering specifically addresses the need for high availability for consolidated VHD deployments.

Co-location of VHD devices onto a single storage LUN should also be done with consideration given to addressing the cumulative workload. Specifically in cases where an application such as Microsoft SQL Server or Microsoft Exchange Server is deployed within a virtual machine, best practices should be applied to ensure that the underlying storage is able to support the anticipated workload. When co-located VHD devices are placed on a common storage volume (LUN), this device should be provisioned to ensure that it can satisfy the cumulative workload of all applications and operating systems located on the VHDs. In cases where storage has been underprovisioned from a performance perspective, all co-located applications and virtual machines may be adversely affected.

Availability and mobility for virtual machines

After initial deployment of a virtualized infrastructure, there is often a need to provide high availability for the services running within the environment. In this case it is often necessary to consider making the virtual machines themselves highly available. In its simplest form, it is possible to provide a high-availability solution by configuring multiple Windows Server 2008 servers into a Failover Clustering configuration. This style of configuration can support up to 16 physical servers that are clustered at the parent partition level. The virtual machines configured on shared SAN storage then become resources that can be moved amongst the nodes.

Windows Failover Clustering – parent partition

The implementation of Windows Failover Clustering for use with Hyper-V virtual machines is identical in its implementation as for any other Windows Cluster environment for other applications such as SQL Server or Exchange Server. Effectively, virtual machines become another form of application that Failover Clustering is able to manage and protect.

To convert an existing virtual server instance into a highly available configuration, the Failover Cluster management is used to configure a new application. Within the wizard, the option to configure a virtual machine needs to be selected as shown in Figure 18. It will be necessary for the virtual machine to be shut down to configure it for high availability, and it will also be necessary for all storage objects, including items such as ISO images that are mounted to the virtual machine, to be located on SAN storage.

Failover Clustering assumes that access to storage objects from all nodes within the cluster are symmetrical. This assumes that all drive mappings, file locations, and mount points will be identical, and during configuration, checks will be made to ensure this condition is met. Warnings will be provided when Failover Clustering is not able to verify some of these aspects. Wizard failure will result when mandatory requirements are not met. In all cases, additional information regarding the warning or failure will be provided to allow for corrective action to be taken.

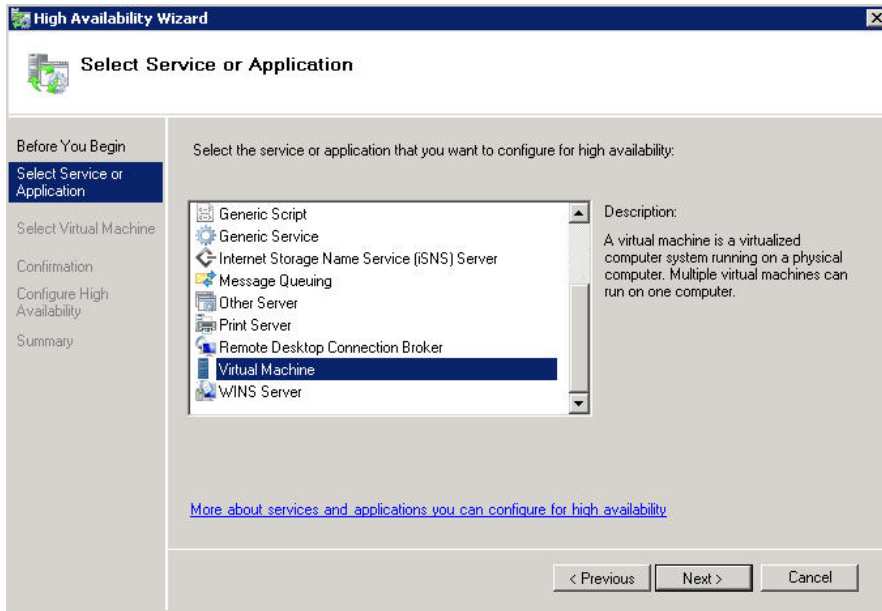


Figure 18. High Availability Wizard

Once imported into Failover Clustering, the virtual machine should be managed and maintained through the Failover Cluster management interface. Specifically, starting and stopping of the virtual machine should not occur outside of the control of Failover Clustering. Should the virtual machine be shut down outside of the control of Failover Clustering, the clustering software will assume that the virtual machine has failed and will restart the virtual machine.

Failover Cluster manager, where necessary, will launch the required virtual machine management interfaces. All availability options and state changes for the virtual machine should be managed through Failover Clustering. Figure 19 details the view of a virtual machine that has been imported via the High Availability Wizard. In this instance, the virtual machine utilizes a VHD located on one of the cluster disk devices, and the remaining two cluster disk devices are utilized as pass-thru devices by the virtual machine.

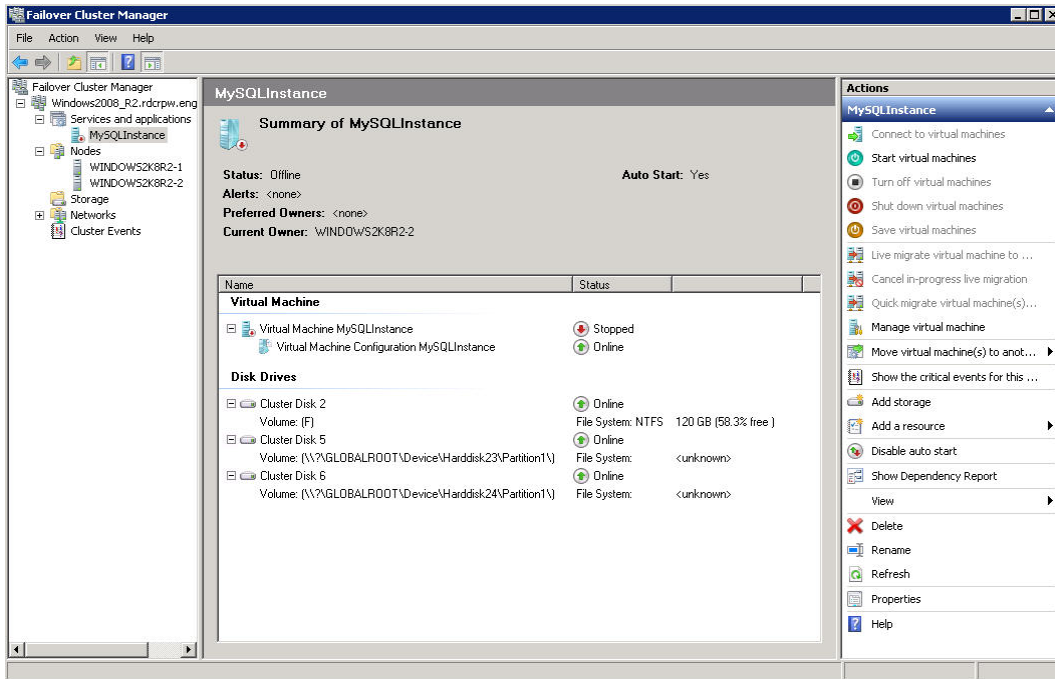


Figure 19. Virtual machine configured with VHD and pass-thru disk devices

A virtual machine instance imported into a high-availability configuration will need to include all related storage disk devices such that the virtual machine can be managed correctly. The High Availability Wizard will fail if it is unable to include all storage configured for the virtual machine within the cluster environment. Therefore it is important to ensure that all shared storage is appropriately configured across all cluster nodes. Subsequent addition of disk storage devices will require that the new storage is also configured appropriately as shared storage within the cluster.

Windows Server 2008 Failover Clustering has a primary goal of maintaining availability of the virtual machine in those cases where the virtual machine becomes unavailable due to unforeseen failures. However, this protection does not always infer that the virtual machine state will be maintained through such transitions. As an example of this style of protection, consider the case of a physical node failure where one or more virtual machines were running, Windows Failover Clustering will detect that the virtual machines are not operational and that a node is no longer available and will attempt to restart the virtual machines on a remaining node within the cluster configuration. Such operations require a restart process and will cause any running applications to be completely restarted.

Movement of virtual machines within a cluster, that is, a proactive move request either by an administrator or by some automated management tool, will cause the virtual machine state to be saved to disk, and then resumed once disk resources have been moved to the target node. This Move operation may also be referred to as a “Quick Migrate” operation. Such move requests will cause an outage from a client application due to the length of time such a move request may take, although the virtual machine state will be resumed. These proactive requests allow Failover Clustering mechanisms to invoke those processes available to coordinate and protect the state of the virtual machine.

Windows Failover Clustering – child partition

Providing high availability of virtual machines through the use of Windows Failover Clustering at the parent level can provide availability for the virtual machine resources. However, protection at the virtual machine level may not be able to provide high availability for the application running within the virtual machines. Consider the case where a virtual machine instance may have suffered from a logical corruption of files that renders the server instance unable to start. The high-availability protection for the virtual

machine may ensure that the virtual machine is in a running state but is unable to ensure that the operating system itself, or the application installed in the server is accessible.

Windows Failover Clustering provides support for application level checking to ensure that services are accessible. For example, a clustered SQL Server instance continually undergoes “Look Alive” and “Is Alive” checks to ensure that the SQL Server instance is accessible to user connections. Implementing clustering within the child virtual machines can provide this additional level of protection.

It is not possible to implement a Failover Cluster configuration within virtual machines running Windows Server 2008 or Windows Server 2008 R2, as a result of the filtering of the necessary SCSI-3 Persistent Reservation commands. However, it is possible to form Windows Cluster configurations with virtual machines running Windows Server 2008 and Windows Server 2008 R2 with the usage of iSCSI shared storage devices. In such configurations, the iSCSI initiator is implemented within the child virtual machines, and the shared storage is defined on the iSCSI LUNs.

Virtual machine Live Migration

When the operating system implemented on Failover Cluster nodes (parent partitions) is Windows Server 2008 R2, it is possible to implement the new Live Migration functionality available with the clustering environment. Live Migrations are a proactive process executed to move virtual machines transparently between nodes. Unlike Move requests, there is no outage from a client application, and the migration between nodes is completely transparent. To achieve this level of client transparency, Live Migrations copy the memory state representing the virtual machine from one server to another so as to mitigate any loss of service. However it should be noted that high-availability protection for virtual machines in the event of failures, such as a physical node failure, will continue to be restart solutions

Live Migration configurations will require a robust network configuration between the nodes within the cluster so as to optimize the memory copy between the nodes to enable the virtual machine transition to be efficient. It will generally be required that such Live Migration configurations will implement at least one dedicated 1 Gb (or greater) network between cluster nodes to enable the memory copy.

When a Live Migration is executed, Failover Clustering will begin a process to replicate the virtual machine configuration and memory state to the target node of the migration. Multiple cycles of replicating the memory state will begin to take place, in an effort to reduce the amount of changes that need to be replicated on subsequent cycles of memory replication. The execution of this memory replication process can be seen through the Failover Cluster Manager console, as shown in Figure 20. Given the ability of the network connectivity to allow for the timely transfer of state, the migration process will, as a final phase, momentarily suspend the machine instance, and switch all disk resources to the target node. After this process, the virtual machine will immediately resume processing. The transition of the virtual machine is required to complete within a TCP/IP timeout interval such that no loss of connectivity is experienced by client applications.

The Live Migration process is different to the Move operation as no suspension of virtual machine state to disk is conducted. Failover Clustering will still provide support for Move operations to be executed. The Move operation may also be referred to as a “Quick Migrate.”

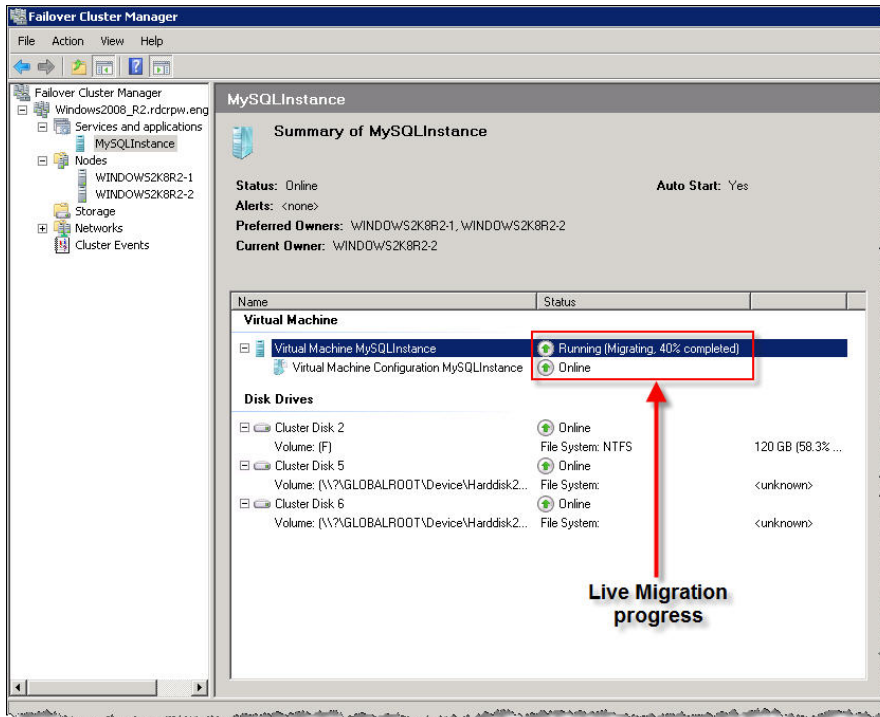


Figure 20. Cluster Live Migration progress

In the event that the migration of the virtual machine will not be able to execute successfully, the migration process will revert the virtual machine back to the originating node. This process will also maintain the availability of the virtual machine to ensure that client access is not impacted. It is also possible to terminate a Live Migration by utilizing the “Cancel in progress Live Migration” option from within the Cluster Manager console.

Windows Failover Clustering Cluster Shared Volumes

Windows Server 2008 R2 also introduces the ability to configure shared SAN storage volumes such that all nodes within a given cluster configuration are able to access the volume concurrently. In this configuration, the volume is mounted as read/write to all nodes at the same time, which differs to the previous mechanism that made the volume read/write to only that node that was the owner of the resource group in which the disk object was configured. The new model for allowing direct read/write access from multiple cluster nodes is called Cluster Shared Volumes (CSVs), and is explicitly implemented to support running multiple virtual machines on different nodes where the VHD storage devices are located on a commonly accessible storage device.

For implementations where there is a requirement to implement storage devices as pass-thru disk objects as previously described, the CSV functionality will not be applicable. However, as previously documented, functionality such as Live Migration will still be available to those configurations. Live Migration of virtual machines is not dependent on CSVs.

CSVs do, however, make the transition process for VHD ownership much more efficient as no transition of ownership and subsequent mounting is required, as is typical of cluster storage devices. The SAN storage configured as CSVs is mounted and accessible by all cluster nodes. To enable the CSV functionality, select the Enable Cluster Shared Volumes context menu, or select the **Enable Cluster Shared Volumes** option from within Failover Cluster Manager on a Windows Server 2008 R2 cluster, as shown in Figure 21.

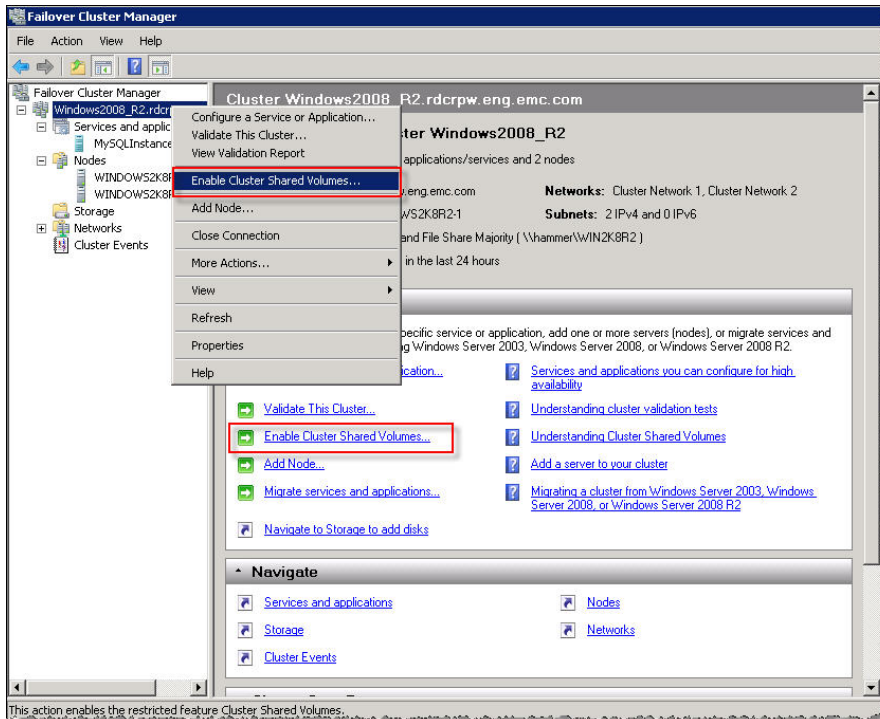


Figure 21. Cluster Shared Volumes from Failover Cluster Manager

As a part of the enabling of the CSV functionality, a warning will be presented with respect to the applicability of the functionality, and providing guidance around support for deployments with CSV operations. In effect, CSV volumes may only be used to store VHD files for virtual machines.




Figure 22. Support notification for CSV functionality

Once enabled, a new Cluster Shared Volumes option will be presented within Failover Cluster Manager, as shown in Figure 23. This selection will allow for the creation and ongoing management of CSV devices. Storage previously configured as cluster storage devices can now be converted to CSV devices, although only devices that have not been allocated for use in other resource groups will be available as devices that can be added. Therefore, to add existing storage devices that host virtual machines, it will be necessary to deconfigure the virtual machines, remove the storage device from any resource groups, convert the storage to CSV, and then reconfigure the virtual machines. The SAN device to be added to the CSV configuration must also contain at least one NTFS volume, as raw disk storage (without NTFS) is not supported.

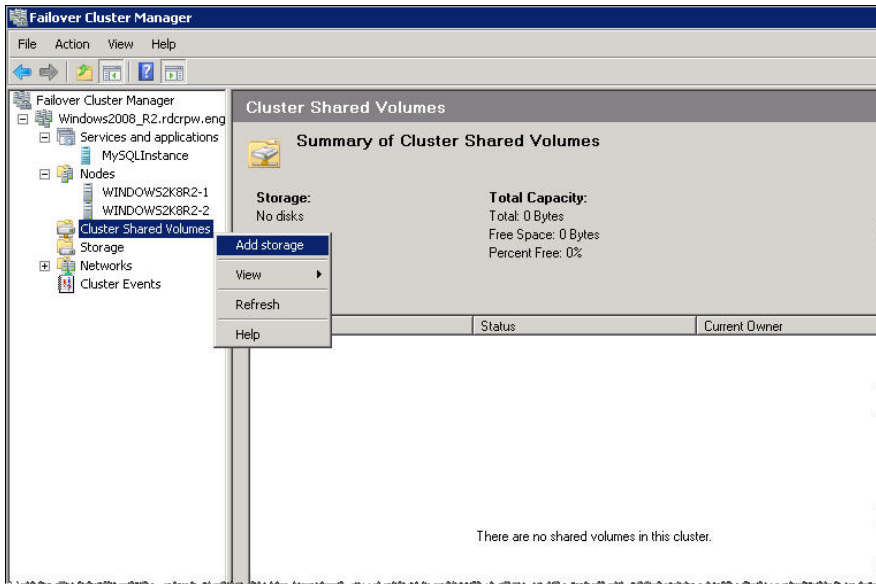


Figure 23. Cluster Shared Volume management

After converting a SAN device to be used as a CSV volume, the storage device will be accessible on all cluster nodes. Additionally the CSV volume will be mounted to a common (but local) location on all nodes, ensuring that the namespace to VHD objects will be identical on all cluster nodes. Figure 24 demonstrates two CSV volumes configured within a Windows Server 2008 R2 environment. The namespace attributed to each CSV volume is based on the System Drive location (which needs to be the same for all cluster nodes) followed by a “ClusterStorage” location, below which the volumes will be physically mounted on each node. The mount location itself will be a sequentially generated name of the form “Volume1” where the appended numeric value is incremented for each subsequent volume.

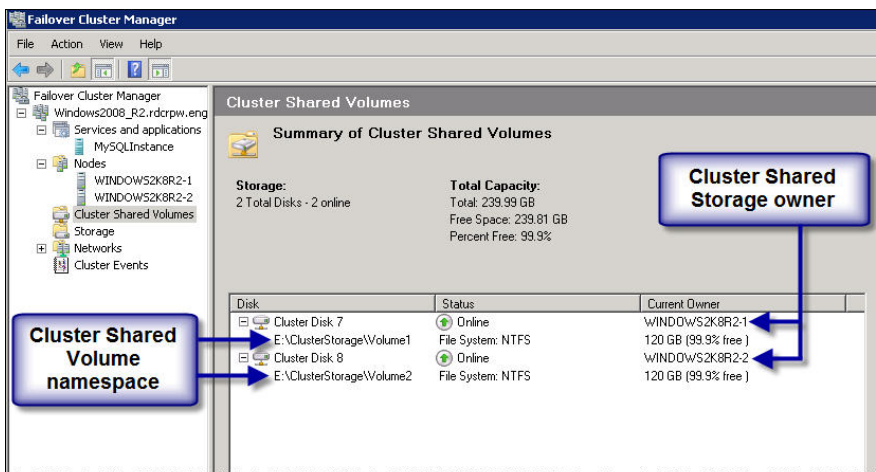


Figure 24. Cluster Shared Volumes added to a cluster

All CSV devices will also indicate a current owner for the resource. The owner is ultimately responsible for coordinating access to the various VHD devices that represent virtual machine storage within the cluster. Virtual machines themselves will continue to run on only a single physical server at any time. When a virtual machine deployed on CSV storage configured within the cluster is to be brought online, the node that is starting the virtual machine will communicate with the CSV owner to register its intent to start generating I/O to the VHD device as the virtual machine is brought into operation. In effect, the node starting the virtual machine will lock the VHD device to ensure no other process is able to write to the VHD from any other node. If the VHD has already been locked by another node, then the request will be

denied. When granted, the node will generate direct I/O to the VHD on the storage device as necessitated by the virtual machine.

CSVs also include a level of protection against external failure scenarios, such as physical connectivity loss from a given node. In the event that connectivity from a node is lost to the underlying storage, I/O operations will be seen to fail, which will invoke a redirection of the I/O over the network to the current owning node. This functionality prevents the abnormal termination of a virtual machine as a result of the loss of storage connectivity. While this functionality will allow the virtual machine to survive from this typically fatal situation, this indirection should not be relied upon to provide ongoing access to the virtual machine. Performance when running in redirected mode may be adversely affected, and immediate attention to resolve the loss of connectivity, or executing a Live Migration, should be undertaken.

Sizing of Cluster Shared Volumes

Windows Server 2008 R2 Cluster Shared Volumes are ostensibly NTFS volumes, and are therefore subject to the same limits of NTFS. NTFS volumes and therefore CSVs have a theoretical maximum of the largest NTFS volume, which is 256 TB (http://www.microsoft.com/whdc/device/storage/GPT_FAQ.mspx). Administrators will need to determine appropriate sizing for CSV volumes based on the cumulative workload expected from the VHD files located within the CSV.

The Cluster Shared Volume will be physically represented by a single LUN presented from a Symmetrix array. The LUN itself will be comprised of one or more hypervolumes within the array, which are in turn allocated on physical disks. Typical sizing for both storage allocation and I/O capacity should be undertaken to ensure that both the storage allocation for a given CSV and the I/O requirements are adequately met.

Undersizing the LUN for I/O load will result in poor performance for all VHDs located on the CSV, and subsequently for all applications installed in the virtual machines utilizing the VHDs. EMC recommends taking a balanced approach to sizing the CSV devices, and adding multiple CSVs as an approach to distribute workloads across available resources.

Site disaster protection with Windows Geographically Dispersed Clustering

EMC has developed the Cluster Enabler product to allow for seamless integration of multi-site storage replication into the framework provided by Windows Failover Clustering. Microsoft Windows Server has supported storage-based multi-site replication scenarios under the Geographically Dispersed Cluster validation program. Compatible solutions may be found within the Windows Server Catalog at <http://www.windowsservercatalog.com/>. The EMC Cluster Enabler product (formally called GeoSpan) has been a supported product for several years under the Windows Geographically Dispersed Clustering program under the product name of EMC SRDF[®]/Cluster Enabler (SRDF/CE).

The implementation of the Cluster Enabler product with EMC Symmetrix Remote Data Facility (SRDF) provides full support of Failover Cluster configurations with multiple forms of storage-based replication, including SRDF/Synchronous and SRDF/Asynchronous, across sites. Due to the tight integration with the Windows Failover Cluster framework, valid supported Failover Cluster configurations and deployed applications are fully supported under the SRDF/CE solution set. This includes support of Windows Hyper-V virtual machines within the SRDF/CE product. Windows Server 2008 R2 Hyper-V configurations utilizing the Windows Server Core installation mode are fully supported with SRDF/CE.

The steps for installation of the SRDF/CE product within a Microsoft Cluster environment are beyond the scope of this paper and are covered in detail in the *EMC SRDF/Cluster Enabler Product Guide*.

When installed, SRDF/CE is transparent to the operations of the typical Failover Cluster Management framework. Installation of the products will install a small number of services and a cluster resource. In Figure 25, the cluster resource for SRDF/CE is indicated. The name assigned to the resource will be preceded with “EMC_” and append the name of the specific resource group.

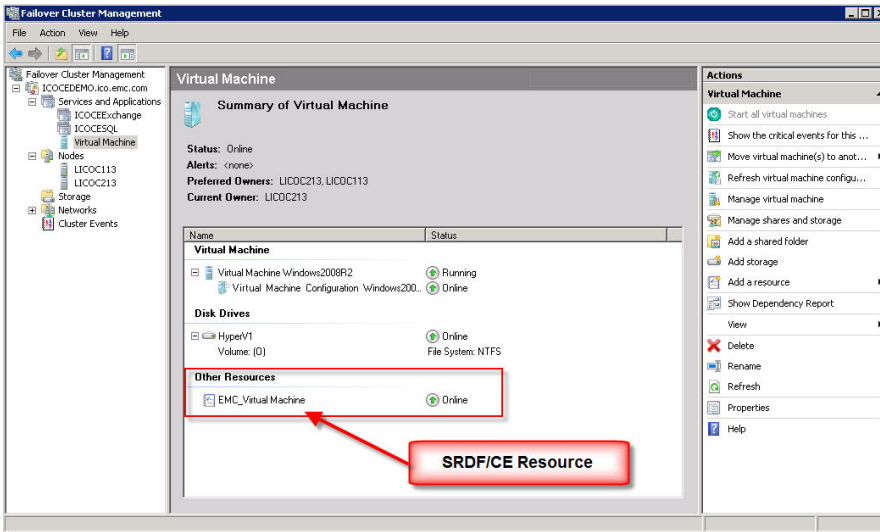


Figure 25. Failover Cluster Manager with an SRDF/CE resource

The clustered disks within the resource group have their dependencies modified to include the SRDF/CE resource defined for the group. This ensures that transitions to other nodes are coordinated appropriately. For “lateral” movements, that is, movement to nodes that are within the same site as the owning node, no transition of SRDF state is required. If the resource is requested to be moved to a “peer” node, that is, a node that is located in the remote site, the SRDF/CE resource will coordinate with other coordination services to transition to the remote disk to a read/write state. Management of the SRDF state of disk devices is fully managed by the SRDF/CE resource, and this functionality is transparent to the Administrator.

The SRDF/CE environment includes a management console, shown in Figure 26, to configure and manage SRDF/CE specific functionality. All typical Failover Cluster functionality is managed through the Failover Cluster Management Console.

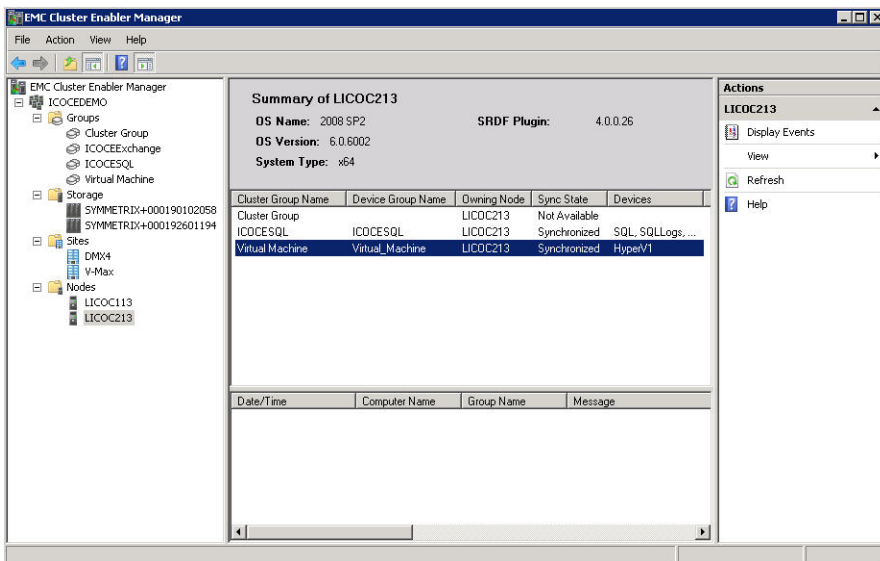


Figure 26. SRDF/CE Management Console

The SRDF/CE Management Console can be used to identify resources used within the various groups configured in the geographically dispersed cluster environment. The management framework utilizes a logical construct of sites, and logically displays resources based on this layout. Figure 27 demonstrates the view of a site, which includes the Symmetrix array located within the site, the cluster nodes (in this case

only the single node LICOC213 is located in the site) and the resources currently located within the nodes in the site.

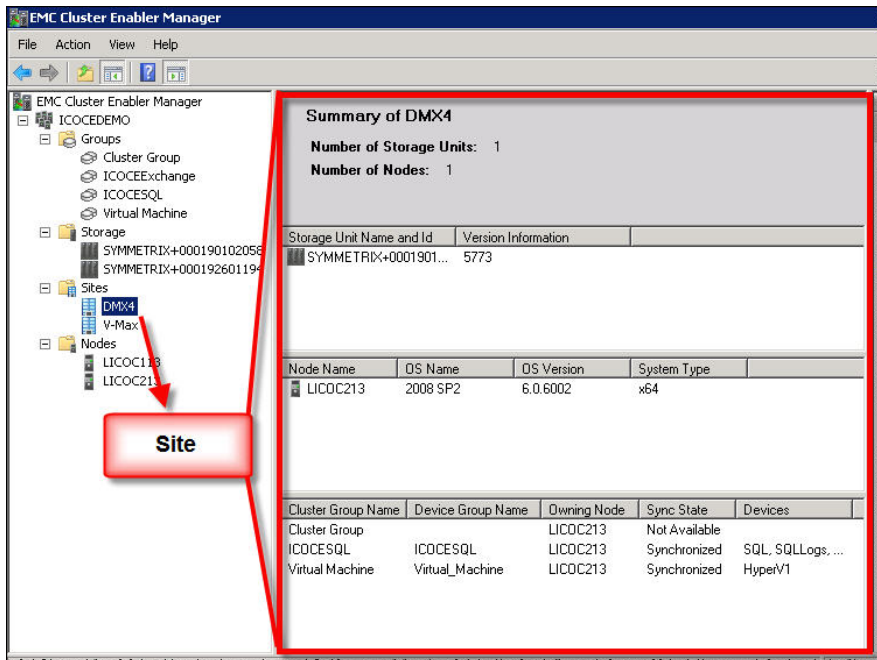


Figure 27. SRDF/CE Site view

All movement, configuration of resources, and online/offline status changes of resources within the cluster will continue to be executed through the standard Failover Cluster Management Console. The SRDF/CE Management Console is used to configure newly created resource group integration with the SRDF/CE framework, or to introduce new shared disk resources into the cluster configuration.

Due to the transparent implementation of the SRDF/CE product, all configurations supported by Windows Failover Clustering are generally supported by SRDF/CE. A notable exception for the current release of SRDF/CE is no support for Clustered Shared Volumes. The design goal for CSV functionality assumes equal access to the CSV volumes from all nodes, whereas the SRDF functionality limits access to the target storage devices in an SRDF relationship. Clustered solutions for Hyper-V virtual machines that do not utilize CSV are fully supported. A subsequent release of SRDF/Cluster Enabler will provide support for CSV functionality.

Details on the configuration and management of the SRDF/CE product may be found in the *EMC SRDF/Cluster Enabler Product Guide*.

Microsoft System Center Virtual Machine Manager

The Microsoft System Center Virtual Machine Manager (SCVMM) product allows for the efficient management of a Hyper-V environment that may incorporate hundreds of physical servers. SCVMM integrates with the various availability products such as Failover Clustering to provide a centralized management, reporting, and alerting platform. SCVMM also has the capability to provide management services for VMware servers and their virtual machine resources.

The centralized management console, as shown in Figure 28, allows administrators to have a centralized view of all managed servers and resources. From with the management console, existing virtual machines can be discovered, deployed, or migrated between managed physical servers. This functionality allows for a dynamic approach to managing physical and virtual resources within the landscape, and to adapt to changing business demands.

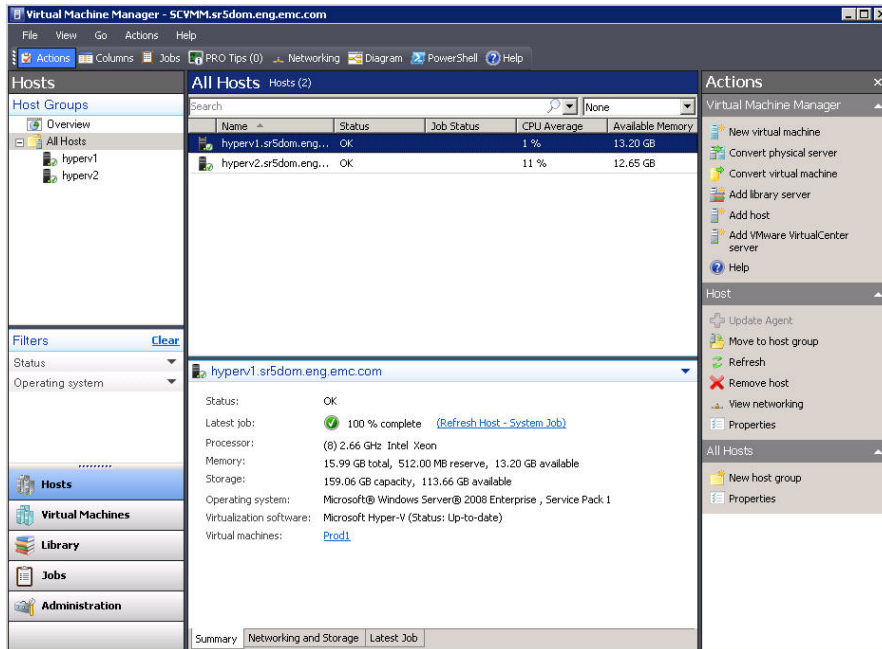


Figure 28. System Center Management Console

The SCVMM product implements a solution for migration of virtual machines between physical servers that does not rely on Windows Failover Clustering functionality as described in the “Windows Failover Clustering” section. Rather, SCVMM is capable of initiating a migration that includes the physical movement of VHD resources, and configuration details of the virtual machine between physical servers. This functionality is implemented in a number of distinct forms. In the first instance, virtual machines, their VHD, and configuration files may be copied across network resources from a source parent partition to a target parent partition. Such network-based file transfers can require considerable network resources and take a large amount of time. A solution based on this functionality does not require any specific storage functionality, and is fully supported on a Symmetrix platform.

SCVMM is also able to initiate SAN-based mechanisms to enable the transfer of resources between parent partitions. This SAN-based migration is implemented in two distinct manners – the first is implemented through the usage of N-Port ID Virtualization (NPIV) as provided by Fibre host bus adapters, and alternatively, SCVMM is able to implement this functionality utilizing the Microsoft Virtual Disk Service (VDS) and a compliant hardware implementation.

NPIV functionality is provided through the appropriate host bus adapter driver and SAN Fabric infrastructure. SCVMM effectively integrates with the driver software to register and or deregister N-Port addresses between the parent partitions. As such, SAN storage devices that are assigned to the N-Port are migrated between the various physical servers as the N-Port is moved. The specific requirements of this functionality are beyond the scope of this paper. Compliant configurations need to be qualified through the EMC Support Matrix located at <http://elabnavigator.emc.com/> (restricted to customers).

EMC Symmetrix provides support for the solution implemented through the VDS implemented with the use of the EMC VDS Provider. The provider is an extension to the Solutions Enabler product set that implements the Microsoft Windows Server compliant functionality to allow for SAN storage management. SCVMM can utilize this functionality to implement SAN storage discovery, mapping, and management.

To be able to utilize the SAN migration capabilities of SCVMM with the EMC VDS Provider, it will be necessary to deploy the VDS Provider on the host that is the SCVMM management host, as shown in Figure 29. The SCVMM server itself will also require connectivity to the Symmetrix storage array, so as to be able to provide discovery and management functionality. It is not necessary to install the EMC VDS Provider on any of the parent partitions. SCVMM will be able to uniquely identify any and all devices for a

given virtual machine, and coordinate the movement of that unique virtual machine within the environment.




Figure 29. System Center Virtual Machine Management VDS migration overview

Deployed in this manner, SCVMM will be able to identify all SAN storage devices that are mapped to a given virtual machine instance on any of the managed parent partitions. When a migration is executed, SCVMM will:

1. Save the virtual machine state, if it is running, to disk
2. Deregister the virtual machine from the source parent partition
3. Execute necessary VDS services on the source parent partition to facilitate unmounting of the volumes
4. Execute the necessary VDS calls to facilitate masking operations within the array, via the VDS Provider, to remove the necessary storage devices from the source parent partition
5. Execute the necessary VDS calls to facilitate masking operations within the array, via the VDS Provider, to add the necessary storage devices to the target parent partition
6. Register the virtual machine with the target parent partition
7. Resume the virtual machine from its saved state, if necessary

To execute the SAN-based Move operation, the virtual machine must have all storage devices located on the Symmetrix array that is accessible from the SCVMM server with the EMC VDS Provider. Selecting the virtual machine to be migrated, the administrator may then select the **Migrate** option from the context menu after selecting the virtual machine, or from the Actions area for the virtual machine, as indicated in Figure 30.

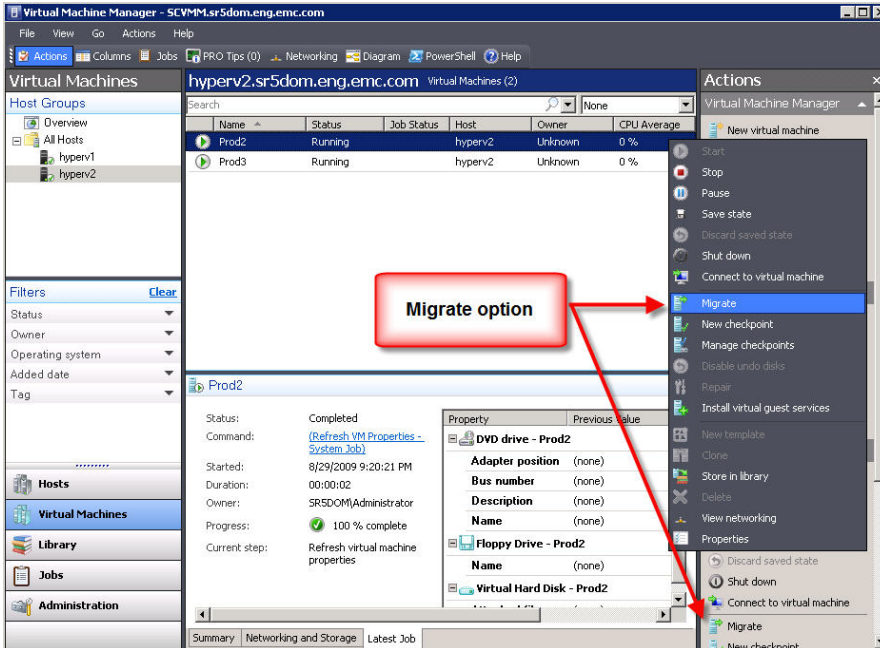


Figure 30. Initiate a virtual machine migration with an SCVMM / VDS mechanism

Selection of the Migrate option will initiate a wizard process that will prompt for various options regarding the transition of state for a running virtual machine, the target parent partition for the migration, network connectivity, and so on. In the Select Host dialog box, it will be possible to verify that the SAN-based migration will be executed and is identified in Figure 31. In the event that storage devices are not located on the SAN device, this selection may appear as “Network”.

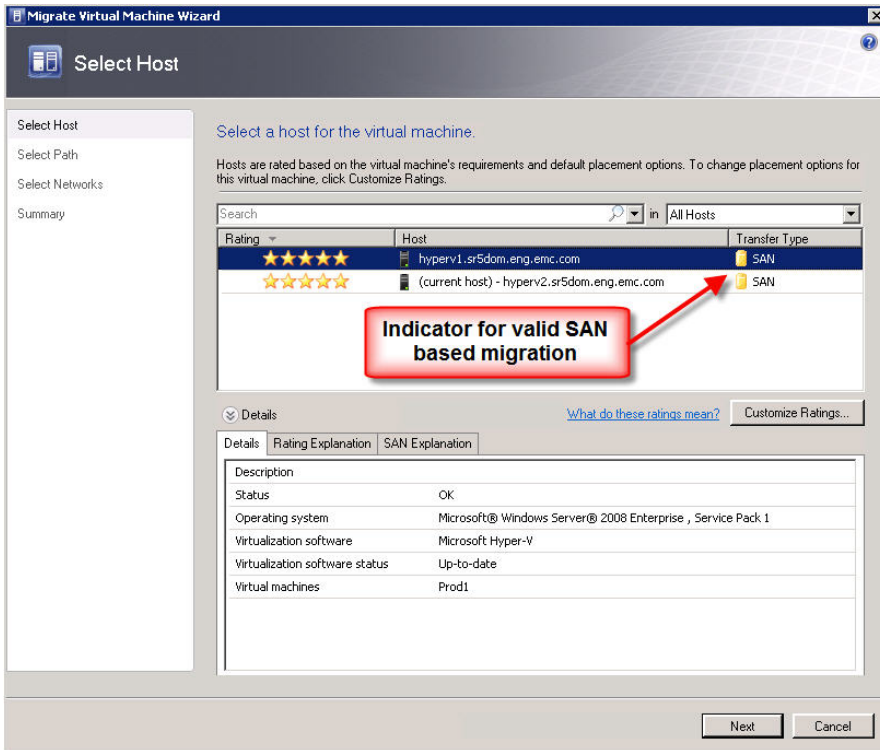


Figure 31. Host and migration style selection

Once the Migrate Virtual Machine Wizard actions have been completed, SCVMM will execute the high-level steps outlined earlier; these steps will be executed as a job, and display to the administrator as shown in Figure 32.

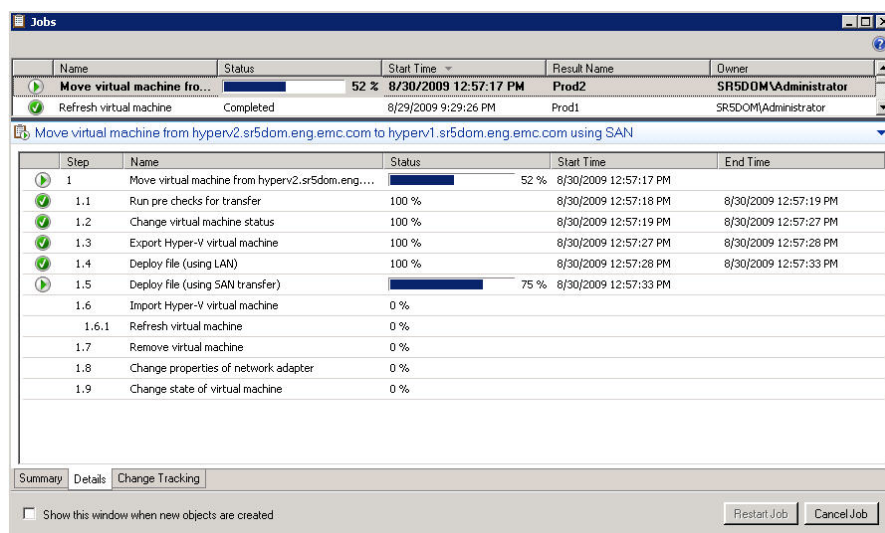


Figure 32. SCVMM virtual machine migrate job steps

The SCVMM implementation of SAN-based migration is a very efficient and powerful mechanism to manage the distribution of virtual machine resources. As the SCVMM product integrates with other System Center products such as Operations Manager, SCVMM may be extended through the Performance and Resource Optimization (PRO) system to automatically initiate such movements in response to changing demands in the distributed environment.

EMC Symmetrix provides support for all System Center Virtual Machine Manager operations, which utilize storage management features through the EMC Virtual Disk Service Hardware Provider. Additionally, all layered functionality implemented by SCVMM, such as integration with Failover Clustering and virtual machine placement, is fully supported by EMC Symmetrix storage platforms.

Integration with EMC layered technologies

Symmetrix storage systems provide significant functionality beyond those features detailed in the preceding sections. Symmetrix V-Max storage arrays include functionality that provides the foundation for creating Dynamic Virtual Data Center configurations.

EMC Solutions Enabler

EMC Solutions Enabler is a prerequisite for many layered product offerings from EMC. Installation of Solutions Enabler at the parent partition level is fully supported and provides the necessary support for configurations such as SRDF/CE, which run at the parent level. Also fully supported are deployments of Solutions Enabler within child partitions that are using iSCSI storage devices.

Solutions Enabler cannot function against storage devices that are VHD devices, even when the VHD devices are located on Symmetrix storage. The underlying LUN configuration for a storage device that is used for VHD placement is not able to be detected from the child partition.

In certain cases, it may be necessary to implement Solutions Enabler within a child virtual machine that is using pass-thru storage devices presented through the parent partition. EMC supports the installation of Solutions Enabler with a child virtual machine using pass-thru storage devices only when the parent partition is running Windows Server 2008 R2, and when the appropriate settings for the virtual machine have been made.

The EMC Solutions Enabler product implements the usage of extended SCSI commands, which are by default, filtered by the parent partition. A bypass of this filtering is provided with Windows Server 2008 R2 Hyper-V, and this pass-through must be enabled to allow for appropriate discovery options from the virtual machine. Microsoft supports allowing full pass-through of SCSI commands as referenced at [http://technet.microsoft.com/en-us/library/dd183729\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/dd183729(WS.10).aspx). EMC recommends allowing SCSI command pass-through only for those virtual machines where it is necessary.

To disable the filtering of SCSI commands, the administrator can execute the following PowerShell script on a Hyper-V parent partition. In the example provided, the name of the virtual machine being auctioned is passed to the PowerShell script when it is executed.

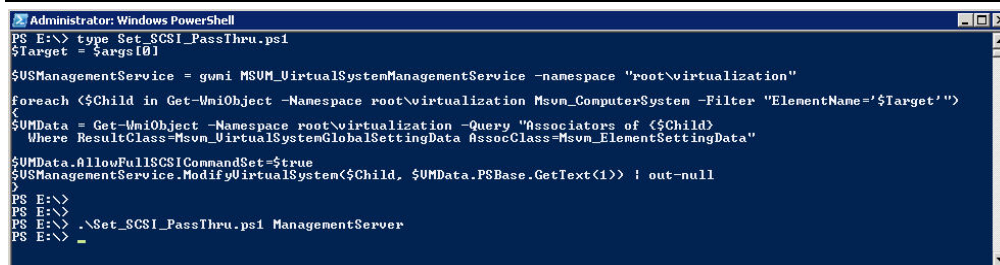
```
$Target = $args[0]

$VSMManagementService = gwmi MSVM_VirtualSystemManagementService -namespace
"root\virtualization"

foreach ($Child in Get-WmiObject -Namespace root\virtualization
Msvm_ComputerSystem -Filter "ElementName='$Target'")
{
    $VMData = Get-WmiObject -Namespace root\virtualization -Query
"Associators of {$Child}"
    Where ResultClass=Msvm_VirtualSystemGlobalSettingData
AssocClass=Msvm_ElementSettingData"
    $VMData.AllowFullSCSICommandSet=$true
    $VSMManagementService.ModifyVirtualSystem($Child, $VMData.PSBase.GetText(1)) |
out-null
}
```

In Figure 33, an example of the execution of the script is provided. In the example, the script is first displayed, and then the virtual machine named “ManagementServer” is provided as the target for disabling SCSI filtering.

The script is provided as-is, and includes no validation or error checking functionality.



```
Administrator: Windows PowerShell
PS E:\> type Set_SCSI_PassThru.ps1
$Target = $args[0]

$VSMManagementService = gwmi MSVM_VirtualSystemManagementService -namespace "root\virtualization"

foreach ($Child in Get-WmiObject -Namespace root\virtualization Msvm_ComputerSystem -Filter "ElementName='$Target'")
{
    $VMData = Get-WmiObject -Namespace root\virtualization -Query "Associators of <Child>"
    Where ResultClass=Msvm_VirtualSystemGlobalSettingData AssocClass=Msvm_ElementSettingData"
    $VMData.AllowFullSCSICommandSet=$true
    $VSMManagementService.ModifyVirtualSystem($Child, $VMData.PSBase.GetText(1)) | out-null
}
PS E:\>
PS E:\>
PS E:\> .\Set_SCSI_PassThru.ps1 ManagementServer
PS E:\> _
```

Figure 33. Example of disabling SCSI filtering on a child partition

It is also possible to check the current value of the SCSI filtering. The following PowerShell scripts may be executed to report on the current SCSI filtering status. Again, it will be necessary to provide the name of the virtual machine target to be reported on.

```
$Target = $args[0]

foreach ($Child in Get-WmiObject -Namespace root\virtualization
Msvm_ComputerSystem -Filter "ElementName='$Target'")
{
    $VMData = Get-WmiObject -Namespace root\virtualization -Query "Associators of
{$Child}"
    Where ResultClass=Msvm_VirtualSystemGlobalSettingData
AssocClass=Msvm_ElementSettingData"
```

```
Write-host "Virtual Machine:" $VMData.ElementName
Write-Host "Currently Bypassing SCSI Filtering:" $VMData.AllowFullSCSICommandSet

}
```

Once set, the setting will persist for the virtual machine, as the setting is recorded against the virtual machine configuration. However it will be necessary to restart the child partition, once the setting has been changed, for the setting to take effect.

Enhanced Virtual LUN Technology

Enhanced Virtual LUN Technology enables transparent, nondisruptive data mobility among storage tiers within the same array and across RAID protection schemes. Enhanced Virtual LUN Technology includes full support for metavolumes.

This functionality allows Hyper-V administrators implement storage tiering for virtual machine infrastructures. Depending on the storage requirements, VHD devices may be migrated to storage tiers that provide the necessary performance and RAID protection as designated by administrators.

Enhanced Virtual LUN Technology offers two types of data movement: migration to unconfigured space and migration to configured space. In each case, the migration provides administrators the ability to move data between high-performance disks and high-capacity disks, or to dynamically populate newly added disk drives.

Hyper-V administrators are able to utilize this functionality in a number of ways to either address inadvertent misplacement of VHD LUNs on underperforming devices, or to provide a mechanism to implement Information Lifecycle Management (ILM) for virtual machines. In the former condition, a system administrator may identify that a certain virtual machine is not performing adequately due to an inappropriate selection of RAID type, or due to placement on physical drives that are suffering from high aggregate workloads. In this case, it is possible to identify devices or free storage areas that may be used to migrate the existing LUN. The migration of the LUN will subsequently migrate the VHD files that reside on that LUN, while providing continuous access to the data and therefore mitigating any loss of availability for the virtual machines. It is important to note that the migration will function on a LUN level, and will affect all the VHD devices located on the target LUN. Pass-thru devices as well as Cluster Shared Volumes may also be migrated utilizing Enhanced Virtual LUN Technology.

Solutions Enabler provides a command line interface (symmmigrate) to define and execute the migration process. In the following example of this migration capability, the migration of a LUN used as the boot device for a virtual machine is undertaken. The source volume is a single RAID 5 device on 1 TB SATA disks, and the target of the migration is a pre-existing RAID 1 device on 300 GB 15k rpm disks. Thus, the migration in this instance is between RAID levels, and between different tiers of storage.

As shown in Figure 34, the source of the migration is the pass-thru device designated as Disk 21 on the parent partition. The pass-thru device has been mapped as the first IDE controller, and the operating system was installed on this device.

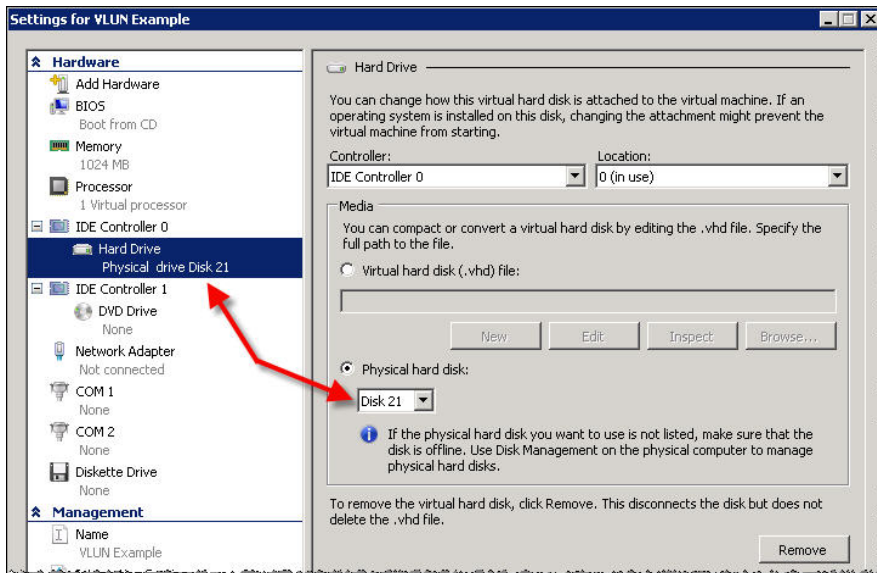


Figure 34. Pass-thru device to be migrated via Enhanced Virtual LUN

The attributes of the LUN can be detailed by using the symprd Solutions Enabler utility from the parent partition. In the example displayed in Figure 35, various device attributes are displayed including the “Device Symmetrix Name” that will be used subsequently in the migration process. Details of the RAID configuration are also shown in Figure 35, and it can be seen that the device is a RAID 5 7+1 configuration.

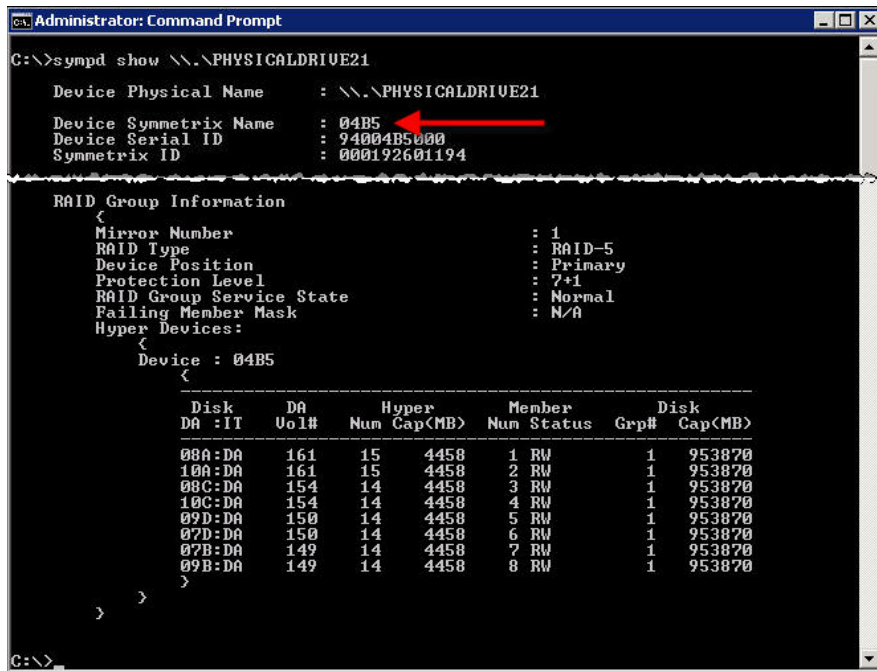


Figure 35. Details of the physical disk device using Solutions Enabler

To process a migration, it is necessary to supply both the source device (4B5) and either a target device to be used to migrate to, or specification of the free space and RAID type to be used. In this example, a target device (454) has already been created as a RAID 1 device, and is of the same size as the source volume. Additional information regarding the options available to process a migration is available in the *Best Practices for Nondisruptive Tiering via EMC Symmetrix Virtual LUN Technical Note*.

Once both source and target devices have been identified, it is only necessary to execute the migration, as shown in Figure 36. In the example, a text file “migrate.txt” contains the source device (4B5) and the target device (454). The subsequent execution of the symmigrate command utilizes the text file and initiates a migration process. A migration is uniquely referenced by the user-defined name provided during the execution of the command. In this instance, the migration is named “VM_Move”.

```

Administrator: Command Prompt
C:\>type migrate.txt
4B5 454
C:\>symmigrate -sid 1194 -file migrate.txt -name UM_Move establish
Execute 'Establish' operation execution for
the device list in device file 'migrate.txt' <y/In> ?y
'Establish' operation execution is in progress for
the device list in device file 'migrate.txt'. Please wait...
Establish Migration.....Started.
Establish Migration.....In Progress.
Establish Migration.....Done.
'Establish' operation successfully executed for
the device list in device file 'migrate.txt'.
C:\>_

```

Figure 36. Execution of a symmigrate operation

The migration will proceed while the virtual machine itself is fully online and accessible to user connections. No outage is incurred during the migration, and is completely transparent to both the parent partition and child partition. It is possible to query the migration process by using the session name provided when the migration was initiated, as shown in first command execution of Figure 37, to determine the current status.

Once the migration has fully completed it will simply be necessary to terminate the migration session, by executing a terminate operation for the session, again utilizing the migration session name. This is shown in the second command execution of Figure 37.

```

Administrator: Command Prompt
C:\>symmigrate -sid 1194 -name UM_Move query -detail
Symmetrix ID: 000192601194
Src  Tgt      Invalid   Status      Done   TGI
Tracks SRC => TGT  (%) Dsk Grp TGI Protection  Session Name
-----
04B5 0454      0 Migrated    100     00 RAID-1
-----
Total
Tracks      0
MB(s)       0.0
Done(%)     100
C:\>symmigrate -sid 1194 -name UM_Move terminate
Execute 'Terminate' operation execution for
the session named 'UM_Move' <y/In> ? y
'Terminate' operation execution is in progress for
the session named 'UM_Move'. Please wait...
Terminate Migration.....Started.
Terminate Migration.....Done.
'Terminate' operation successfully executed for
the session named 'UM_Move'.
C:\>_

```

Figure 37. Completion of the migration process

Once complete, the physical device contents will have been completely migrated to the target device, and will have all the characteristics of the target device. Figure 38 shows the characteristics of the storage device after the migration has been completed. The output is from the execution of the same symcmd command as used in Figure 35.

```

Administrator: Command Prompt
RAID Group Information
<
Mirror Number           : 1
RAID Type               : RAID-1
Device Position        : Primary
Protection Level       : 2
RAID Group Service State : Normal
Failing Member Mask    : N/A
Hyper Devices:
<
Device : 04B5
<
-----
Disk  DA  DA  Hyper  Member  Disk
DA  :IT  Vol#  Num  Cap(MB)  Num  Status  Grp#  Cap(MB)
-----
07C:D3 121  10  31203  1  RW      0  279140
09C:D3 121  10  31203  2  RW      0  279140
>
>
>
C:\>

```

Figure 38. Device attributes after migration has been completed

The ability to dynamically migrate storage devices is also applicable to Cluster Shared Volumes (CSVs) as well as any other storage configuration provided by the Symmetrix V-Max array. Administrators have the capability to apply this powerful, dynamic volume migration capability to implement storage tiering against Hyper-V configurations but targeting both the operating system VHD LUNs as well as those LUNs that are used to host application data.

Auto-provisioning Groups for Hyper-V

Symmetrix V-Max with Engenuity 5874 provides administrators with a simplified model for storage provisioning. This new storage provisioning model is referred to as Auto-provisioning Groups.

Historically, administrators were required to provide somewhat static relationships between storage devices with the Symmetrix array, mappings of those devices to front-end directors for host connectivity, and additionally manage masking operations to ensure that hosts were able to access the requisite storage devices. This methodology served administrators well, and these were often only required to be made once.

Increasingly, administrators need to deal with a dynamic environment, where the introduction of new servers and systems occurs on a regular basis. Deployments of clustered instances of SQL Server are much more commonplace, as is the adoption of technologies such as server virtualization. To assist administrators with the ability to deal with these business challenges, EMC Symmetrix V-Max with Engenuity introduces the new Auto-provisioning Groups functionality, which maps directly into the needs of Hyper-V administrators.

Administrators are now able to define relationships between storage objects and host connectivity, and allow the Engenuity functionality, within the Symmetrix V-Max array, to execute the appropriate changes. This ability to create logical relationships through views also helps to ensure that appropriate devices are automatically included in changes. For example, in a cluster configuration, only a single pool of storage devices needs to be defined. Views created based on this pool of devices ensures that any hosts included in those views will be able to access the required devices. This is in contrast to a manual process where administrators may have to manually ensure that mapping and masking entries have been created.

Additionally, it is possible to define storage devices in the same manner to ensure that they are appropriately managed for a virtualized environment. Hosts may use N-Port ID Virtualization (NPIV) to provide a more dynamic method for allocating storage to virtual machines. NPIV initiators are processed the same as for physical Fibre Channel controller addressing.

The follow steps outline the requirements for implementing Auto-provisioning Groups functionality

1. Create the storage group, which defines the specific Symmetrix devices that will be presented to the host. In the example, a storage group named HyperV_devs is created, and the respective devices that

represent this logical grouping are added, in the example the target devices are A7 and 245, and the Symmetrix V-Max array is 1234

```
symaccess -sid 1234 create -name HyperV_devs -type storage A7,245
```

2. Create the director group, which defines the directors to which the devices are to be mapped, and through which the host will be able to access the devices as defined in the storage group. In this case, the group will be called HyperV_ports and will contain director ports 7e:0 and 10e:0.

```
symaccess -sid 1234 create -name HyperV_ports -type port -dirport 7e:0,10e:0
```

3. Create the host initiator groups, which define the WWNs of the HBAs that are used by the host. In this instance, the target parent server has two HBA WWN ports, and these are configured into an initiator group that bears the hostname. In the example, the WWNs are not explicitly called out.

```
symaccess -sid 1234 create -name HyperV_SRV -type initiator -wwn <<HBA1_WWN>>
```

```
symaccess -sid 1234 add -name HyperV_SRV -type initiator -wwn <<HBA2_WWN>>
```

4. As the final step, the view itself is created. This process will execute all mapping and masking operations between the previously defined groups.

```
symaccess -sid 1234 create view -name HyperV_view -storgrp HyperV_devs  
-portgrp HyperV_ports -initgrp HyperV_SRV
```

After execution of the steps, the storage devices will be provisioned to the server via the Symmetrix V-Max array ports. In the event that the storage devices have not been previously mapped to the specified directors, Symmetrix V-Max Auto-provisioning Groups functionality will add the necessary mapping, and provide a LUN ID value to the devices.

Auto-provisioning Groups and Microsoft Windows Failover Clusters

The Auto-provisioning Groups functionality provides significant value when storage administrators are building or deploying new Hyper-V servers to form a Windows Failover Cluster to implement a highly available configuration for virtual machines. If an existing set of storage and director groups has been defined, then it will only be necessary to provide information for any additional port group configuration. With the addition of the new port group defined for the new host, it is simply a matter of defining an additional view to represent the new host connectivity.

Assuming that the steps defined for the initial initiator groups as outlined in the previous section have been completed, the steps required for defining storage to be used by a new server in a Windows Server Failover Cluster are outlined next.

1. Define a host initiator group for the new server, and add the relevant WWNs for the HBAs located in the specific additional server into the host initiator group.

```
symaccess -sid 1234 create -name HyperV_SRV2 -type initiator -wwn <<HBA1_WWN>>
```

```
symaccess -sid 1234 add -name HyperV_SRV2 -type initiator -wwn <<HBA2_WWN>>
```

2. Create a new view, which encompasses the previously created storage and director groups, and includes the new initiator group for the host.

```
symaccess -sid 1194 create view -name HyperV_SRV2_view -storgrp HyperV_devs -  
portgrp HyperV_ports -initgrp HyperV_SRV2
```

The result of this process will be that the additional server will be able to access the same devices that were defined when the storage group was defined. This process will significantly improve administrative processes, and should result in a reduction of errors caused by the requirement of a greater number of more complex operations. In the provided example, the same port group was utilized for both hosts; however, it is equally possible to define additional port groups if that is appropriate in the given configuration.

Connectivity recommendations

Symmetrix V-Max arrays provide a highly flexible host connectivity environment, which allows for the creation of a scalable Hyper-V configuration. It is recommended to configure at least two host bus adapters (HBA) per Hyper-V parent partition (or physical server) with the goal of presenting multiple unique paths to the Symmetrix V-Max system across multiple directors. The benefits of multiple paths extend beyond the performance improvements, and provide for the creation of a high-availability environment when coupled with appropriate switch and Symmetrix V-Max front-end director connectivity.

To provide the highest levels of availability, all single points of failure need to be addressed. While not a regular occurrence, it may be necessary to occasionally perform director maintenance, including memory upgrades. These procedures may require the removal of the director and its associated connectivity from the Symmetrix V-Max system. As a result, each Hyper-V parent partition should have redundant paths to multiple front-end directors. Each Hyper-V parent partition should be connected to even and odd directors within a single V-Max Engine, or across directors within multiple V-Max Engines, when available.

For each HBA port, at least one discrete Symmetrix front-end port should be configured. It is recommended that each HBA port be configured to two Symmetrix front-end ports. Connectivity to the Symmetrix front-end ports should consist of first connecting unique hosts to port 0 of the front-end directors before connecting additional hosts to port 1 of the same director and processor. This methodology for connectivity ensures all front-end directors and processors are utilized, providing maximum potential performance and load balancing for I/O-intensive Hyper-V configurations. This is an important consideration as each virtual machine will typically be providing service for its own application environment.

Figure 39 represents a logical view of a single V-Max Engine and connectivity to two physical Hyper-V parent partitions (physical servers). The configuration implements a highly available and scalable design where Hyper-V hosts are dual-pathed, and each path connects to two separate front-end modules on different directors. Not shown in this graphic is any SAN fabric that should also be configured in a highly available manner.

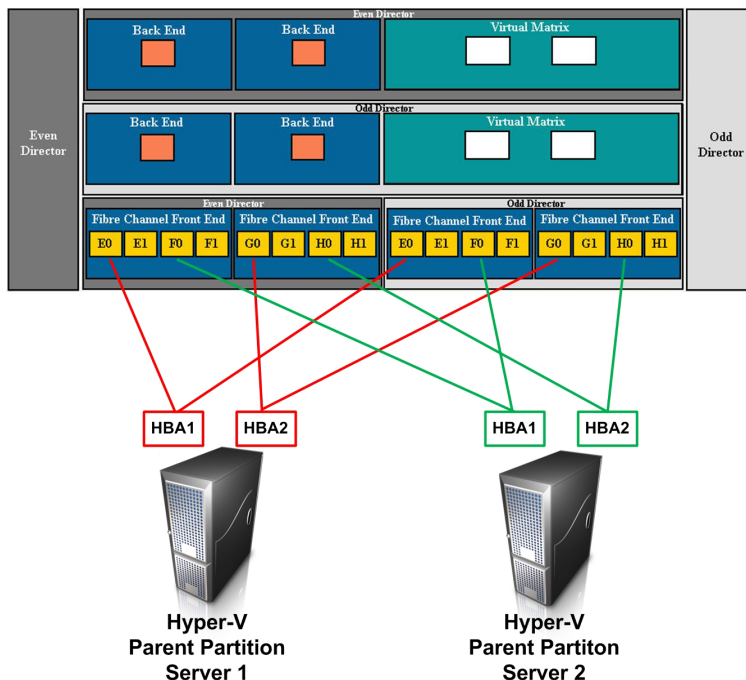


Figure 39. Highly available connectivity for Hyper-V servers

Configurations with multiple paths to storage LUNs will require a path management software solution on the Windows host. The recommended solution is EMC PowerPath[®], which is the industry-leading path management software with benefits including:

- Enhanced path failover and failure recovery logic
- Improved I/O throughput based on advanced algorithms such as the Symmetrix Optimization load balancing and failover “policy”
- Ease of management including a Microsoft Management Console (MMC) GUI snap-in and CLI utilities to control all PowerPath features
- Value-added functionality including Migration Enabler, to aid with online data migration, and LUN encryption utilizing RSA technology.
- Product maturity with proven reliability over years of development and use in the most demanding enterprise environments.

While PowerPath is recommended, an alternative is the use of the Multipath I/O (MPIO) capabilities native to the Windows operating system. The MPIO framework has been available for Windows for many years; however, it was not until the release of Windows Server 2008 where a generic device specific module (DSM) from Microsoft was included to manage Fibre Channel devices. For more information regarding the Windows MPIO DSM implementation, please see the “Multipath I/O Overview” topic at <http://technet.microsoft.com/en-us/library/cc725907.aspx>.

Conclusion

Symmetrix storage arrays provide an extremely scalable storage solution, which provides customers with industry-leading capabilities to deploy, maintain, and protect Windows Hyper-V environments. Additionally, the Symmetrix V-Max architecture implements a new strategy in the Symmetrix scale-out solutions for applications such as Microsoft Windows Hyper-V, providing flexible data protection options to meet different performance, availability, functionality, and economic requirements. The ability to support a wide range of service levels with a single storage infrastructure provides a key building block to implementing Information Lifecycle Management (ILM) by deploying a tiered storage strategy.

These new technologies provide an easier and more reliable way to provision storage in Microsoft Windows Hyper-V environments, while enabling transparent, nondisruptive data mobility between storage tiers for standard next-generation Symmetrix system volumes. Industry-leading multi-site protection through the SRDF/Cluster Enabler solution allows customers to implement a complete end-to-end solution for virtual machine management and protection.