

Dell™ High Availability Solutions Guide for Microsoft® Hyper-V™

Notes and Cautions



NOTE: A NOTE indicates important information that helps you make better use of your computer.



CAUTION: A CAUTION indicates potential damage to hardware or loss of data if instructions are not followed.

Information in this document is subject to change without notice.

© 2008 Dell Inc. All rights reserved.

Reproduction of these materials in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden.

Trademarks used in this text: *Dell*, the *DELL* logo, *EqualLogic*, *EqualLogic Group Manager*, *Modular Disk Storage Manager*, *Navisphere Manager*, *PowerEdge*, and *PowerVault* are trademarks of Dell Inc.; *Microsoft*, *Windows*, *Windows Server* and *Hyper-V* are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

Contents

1	Getting Started with High Availability	7
2	Clustering Implementations for High Availability	9
	Hyper-V Guest Clustering	9
	Hyper-V Host Clustering	12
	Planned and Unplanned Failover (Host Cluster Model Only)	13
	Planned Failover (Quick Migration)	14
	Processor Compatibility Requirements	16
	Unplanned Failover	16
	Comparison of Guest and Host Clustering	19
3	Designing for High Availability	21
	Node and Disk Majority versus Node Majority Cluster Configuration	21
	Storage Configuration	22
	Sharing a LUN for multiple HA VMs	22
	Storage Options for Hyper-V Host Clustering	23

Network Configuration	24
Cluster Network Types	24
Best Practice Recommendations	26
Recommended Configurations	26
Design for iSCSI-Based Configurations	27
Design for Fiber Channel Configurations	29
Design for Direct-Attach SAS Configurations	31
Integration between Hyper-V and Windows Failover Clustering	33
Microsoft KB951308 QFE for Hyper-V Host Clusters	34
4 Implementing High Availability in Hyper-V	35
Implementing Hyper-V Host Clusters	35
Preparing the Systems	35
Configure the Failover Clustering Feature and Hyper-V Role	36
Creating Virtual Machines that will be HA VMs	37
Enabling VMs as HA VMs	39
Migrating HA VMs Across Hosts in the Cluster	40
Hyper-V Host Clusters on Server Core	41
Setting up the Server Core System for Hyper-V and Failover Clustering	41
Setting up Remote Management to Server Core	42
Remote Management from Windows Server 2008 Full install	42
Remote management from Windows Vista	43

5	Administering High Availability for Hyper-V	45
	Changing the configuration of HA VMs	45
	For Configuring Changes that Impact the VM's HW Profile	45
	For Configuring Changes that DO NOT Impact the VM's HW Profile	46
	Adding Passthrough Disks to HA VMs	46
	Sharing LUNs Between HA VMs	47
	HA VM Resource Group Names	48
	Cluster Resource Group for HA VMs Reported as "Partially Online"	48
	Configuring Rules for Associating Hosts to HA VMs.	48
	Setting Host Preferences for HA VMs	49
	Blocking HA VMs from Running on Certain Hosts	49
	Converting an Existing Non-HA VM to a HA VM	49
	Converting an HA VM to a Non-HA VM	50
A	Using SMB/CIFS Shares in Hyper-V	53
B	Using Mount Points in Hyper-V	55

C	Hyper-V Host Clustering Data Template	57
D	References.....	59
	Glossary	61

Getting Started with High Availability

High Availability (HA) solutions are not new to Windows® environments. Microsoft® has offered native support for HA through its Cluster Services in Windows 2000 and Windows Server® 2003, and with the Windows Failover Clustering feature in Windows Server 2008. Failover clustering reduces the possibility that any single point of failure in your virtualization environment will cause your workloads to become unavailable. In a non-virtualized environment, HA implementation was traditionally only done for environments that were hosting critical workloads because the loss of a non-critical service hosted on any single physical server was considered acceptable.

However, in a virtualized environment in which a single physical server will host many virtual machines, the loss of that physical server will result in the loss of service of multiple workloads. To prevent this scenario, implementing HA should be a primary goal of any virtualization deployment regardless of scale or scope, especially in a production environment.

This *Dell High Availability Solutions Guide for Microsoft Hyper-V* introduces different options to implement HA for Hyper-V environments, and provides recommendations and best practices for configuring your Hyper-V HA environments on Dell hardware.

This Solutions Guide assumes that you have a basic understanding of Hyper-V as well as the networking and storage concepts of Hyper-V. If you have not already read them, Dell strongly recommends that you review the following guides to get an in-depth understanding of Hyper-V solutions on Dell hardware:

- *Dell Solutions Overview Guide for Microsoft Hyper-V*
- *Dell Networking Solutions Guide for Microsoft Hyper-V*
- *Dell Storage Solutions Guide for Microsoft Hyper-V*

Clustering Implementations for High Availability

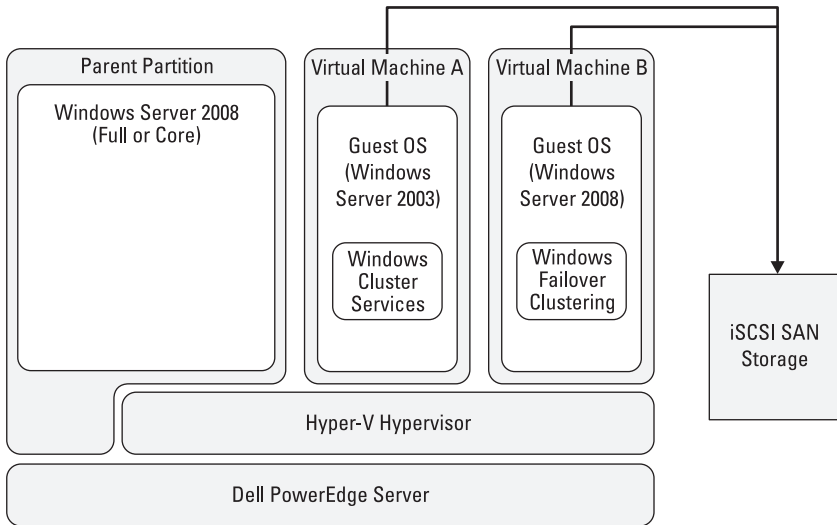
Hyper-V HA implementations can broadly be categorized into Hyper-V Guest Clustering and Hyper-V Host Clustering. The differences in the two types are based on where the clustering service is running and what the clustering service is managing.

Hyper-V Guest Clustering

In a Guest Cluster implementation, the clustering service runs on the guest operating system in the virtual machine. This type of clustering provides high availability for applications that are hosted within the virtual machines.

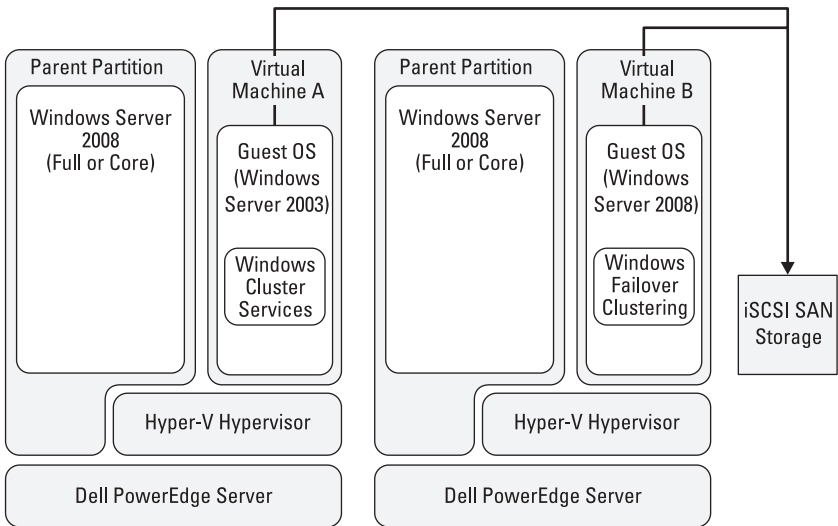
As illustrated in Figure 2-1 and Figure 2-2 below, Hyper-V Guest Clustering can be implemented with the VMs either deployed on a single physical server, or across multiple physical servers.

Figure 2-1. Hyper-V Guest Clustering Within a Single Physical Server



In Figure 2-1, virtual machines A and B are running on the same server. In this scenario, the workloads being managed by the clustering service in the guest operating systems in the VMs will not survive a physical server outage or a crash of the parent partition. If the physical server fails or the Windows Server 2008 instance running in the parent partition crashes, all of the VMs running on this system fail, causing the clustered workloads within the VMs to also fail.

Figure 2-2. Hyper-V Guest Clustering Across Two Standalone Physical Servers



In Figure 2-2 above, the clustered virtual machines A and B are running on different physical servers. In this scenario, the workloads being managed by the clustering software within the VMs are able to offer continued service even in the case of an unplanned physical server downtime or parent partition outage.

As in all cluster implementations, the Hyper-V Guest Cluster implementation also requires that the cluster service running in all the guest OSes have direct access to a common shared storage. In the case of Hyper-V Guest Cluster, the only way to provision shared storage that is accessible directly by multiple VMs is by using iSCSI based storage (Direct Attached SAS and Fiber Channel Arrays are not supported for Hyper-V Guest Clustering). With the iSCSI initiators running in the guest operating systems, Hyper-V VMs can directly talk to the iSCSI storage without going through the storage stack of the parent partition.



NOTE: The *Dell Storage Solutions Guide for Microsoft Hyper-V* provides more detailed information on provisioning iSCSI storage disks directly to Guest OSes.

Because configuring clustering within the guest OSES in the VMs is no different than configuring clustering on physical servers that are not virtualized, this solutions guide does not discuss this model any further. For more information, refer to the documentation for the guest operating system.

Hyper-V Host Clustering

In this implementation, the Hyper-V host servers are clustered by leveraging the Windows Failover Clustering feature in Windows Server 2008 x64 editions. Failover Clustering in Windows Server 2008 is a robust and mature third-generation clustering solution that has been significantly improved in Windows Server 2008. Leveraging this Failover Clustering feature allows Hyper-V to offer HA solutions on a wide array of storage solutions including Direct-Attach SAS, iSCSI, and Fiber Channel storage arrays.

Figure 2-3. Hyper-V Host Clustering

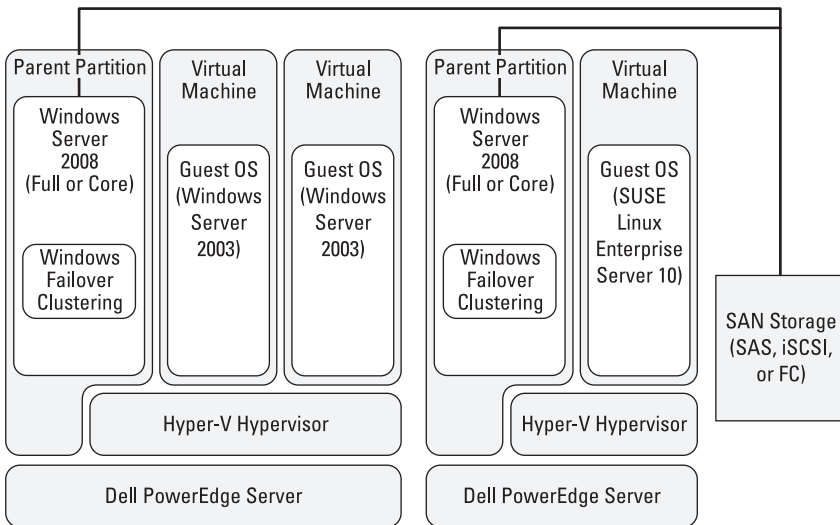


Figure 2-3 illustrates a sample Hyper-V Host Cluster configuration with the Windows native cluster service running in the Windows Server 2008 x64 Parent Partition. In this scenario, the Failover Clustering service in the parent partition will manage the VMs as the cluster resources. The VMs running on either of these servers can be configured as highly available and does not

require clustering capabilities in the guest operating systems in the VMs. The only requirement is that all the virtual machine files are on storage accessible by all the physical servers that are part of the cluster. In this document, we will refer to VMs managed by the cluster for high availability as "HA VMs".

Hyper-V Host Clustering is the recommended mechanism to implement HA in a Hyper-V environment. It has many advantages including a wide array of configuration options, ease of management, and scalability. In addition to providing HA, this solution also supports planned downtime activities such as for physical server maintenance and enabling dynamic datacenter capabilities such as VM workload balancing.

Windows Server 2008 supports a maximum of 16 physical servers in a single Failover Cluster configuration. Therefore a single Hyper-V Host Cluster configuration can have a maximum of 16 physical servers. Hyper-V Host Clustering may be implemented on the following editions of Windows Server 2008:

- Windows Server 2008 x64 Enterprise Edition
- Windows Server 2008 x64 Datacenter Edition

Planned and Unplanned Failover (Host Cluster Model Only)

The Hyper-V Host Cluster model accommodates both planned and unplanned downtime. A planned downtime is when a physical server that is part of a Hyper-V Host Cluster is taken offline by the administrator for maintenance such as a hardware upgrade or applying an OS patch. An unplanned downtime is when a physical server that is part of a Hyper-V Host Cluster becomes offline unexpectedly due to hardware failure, a parent partition crash, or other reason. Physical servers that are part of a cluster are commonly referred to as "cluster nodes".

Planned Failover (Quick Migration)

To accommodate the planned downtime of a server, the administrator can choose to move all the VMs from that server to the other servers in the Hyper-V Host Cluster. Here is a high-level workflow of the VM migration process during planned failover:

- 1 VM in the originating node is placed in save state (the contents of memory are saved to files located on SAN storage)
- 2 SAN storage disk ownership is transferred from the originating node to the destination node
- 3 VM is started on the destination node and state is restored from SAN storage

Figure 2-4. Hyper-V Host Cluster before Migrating

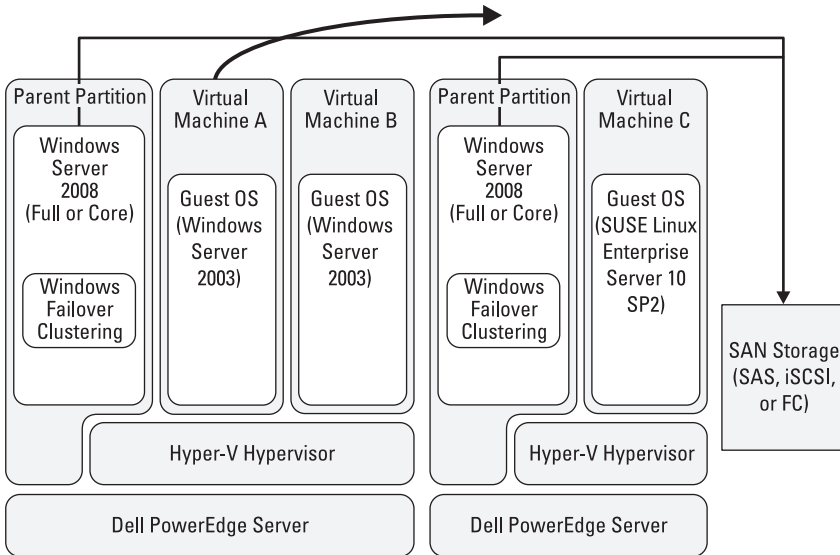
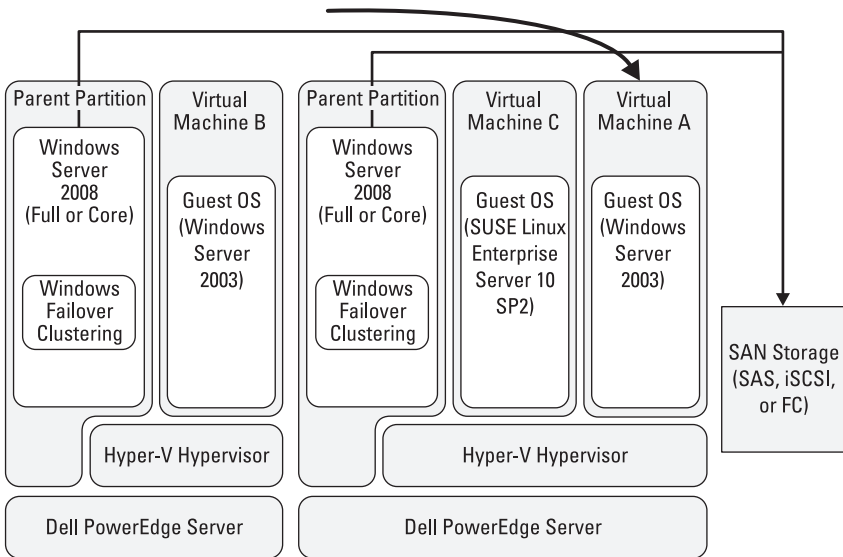



Figure 2-5. Hyper-V Host Cluster after Migrating



This save and restore operation ensures that the state of the VM is retained during the migration process. The guest operating system is unaware of this VM save and restore process: the guest OS is "frozen" during the migration period. Microsoft calls this capability to migrate VMs across physical servers "Quick Migration".

Even though there is no loss of state or data during this process, it is not a "Live" migration process. The guest OS and the workloads within the VM will be offline and not available to users during the migration period.

 **NOTE:** Any active connections with the guest operating system in the VMs will be lost when the VM undergoes migration. The connections have to be re-established after the migration is complete.

The offline duration period is typically between 15 seconds and a few minutes, depending mainly on the following factors:

- **Amount of memory used by the guest operating system:** This is the most important factor because the time taken for VM migration is directly proportional to the amount of memory that is being used by the guest operating system. The more memory that is being used by the guest OS at the time of the migration, the longer the migration time. Note that the

amount of memory being used by the guest OS is not the same as the amount allocated to the virtual machine. Unless the virtual machine is running a very memory intensive workload, the memory utilized is usually much less than the memory allocated to the VM.

- **Speed of the storage fabric:** The connection speed to the storage array is a factor because the state of the VM is transferred from the source node to the shared storage and then restored from shared storage to the destination node over the storage fabric. So it is very important that you plan your network (in case of iSCSI) and storage deployment to ensure that there are no bottlenecks that could impact quick migration times.

Processor Compatibility Requirements

"Quick Migration" is a stateful migration process and so requires that the processors on the source and destination cluster nodes are compatible. Hypervisor-based virtualization software (such as Hyper-V) allows certain CPU instructions from the virtual machines to execute directly on the physical processors. Therefore, the state of a VM may have a dependency on a specific feature of the physical processor.

Processor features vary between processor vendors, processor models, and processor families. During a VM quick migration across physical servers that do not have identical processors, if the processor on the destination node does not support a feature that the VM leveraged in the source node, it can cause the VM state restore process to fail on the destination node.



NOTE: In the scenario when the processors on the cluster nodes are not feature compatible and the VM fails to restore state on the destination node, the cluster service will attempt to restore the VM (along with the saved state) on the originating cluster node if that node is still online. However, this behavior is impacted if changes are made to the default settings for how the Failover Cluster service in Windows Server 2008 manages its resources.



NOTE: Dell strongly recommends that all the nodes of a Hyper-V Host Cluster configuration have identical processors.

Unplanned Failover

In the case of an unplanned downtime, the nodes in the Hyper-V Host Cluster do not have the time to initiate a VM "Quick Migration". Instead, the cluster service automatically restarts all the HA VMs on the other nodes of the cluster. In this unplanned failover scenario, the state of the VM is not

retained and there is possibility for data loss. To the guest operating system, this is similar to an unexpected restart of the operating system. However, the failover of VMs is automatically handled by the cluster service, and workloads running within the VMs are back online in a short period of time. Here is a high-level workflow of the unplanned failover scenario:

- 1 A node in the Hyper-V Host Cluster configuration fails. All states of the VM (including guest OS and workloads within the VMs) that have not been committed to SAN storage are lost.
- 2 Cluster service transfers ownership of SAN storage disk resources from the failed node to other nodes in the Hyper-V Host Cluster.
- 3 VMs are restarted on the other nodes of the Hyper-V Host Cluster.

Figure 2-6. Hyper-V Host Cluster with Both Servers Running

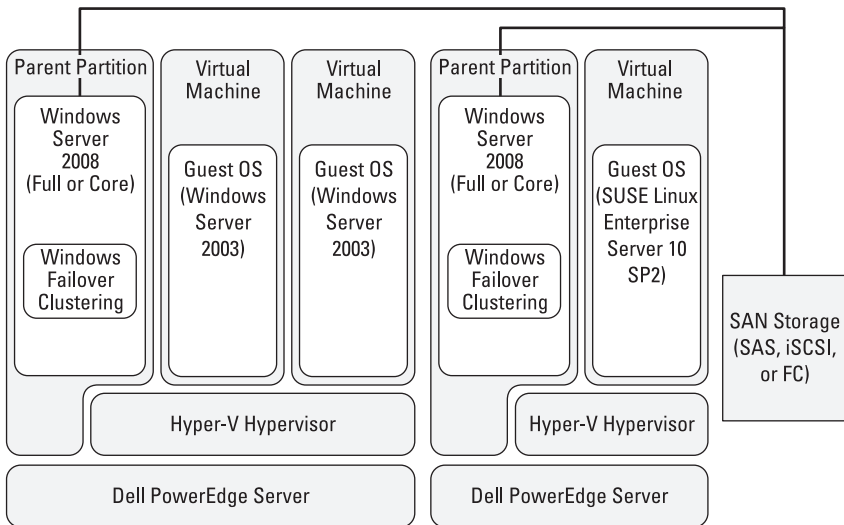
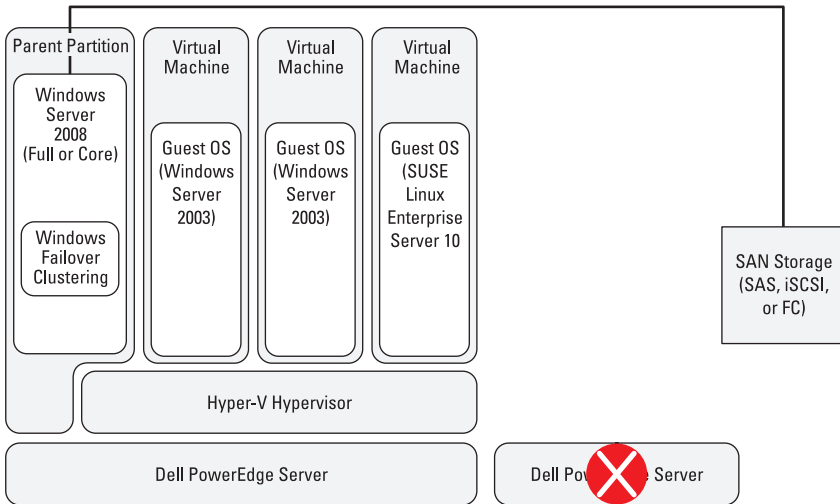




Figure 2-7. Hyper-V Host Cluster after One Server Fails



Unlike in the case of a planned failover where the administrator chooses a specific destination node for the VMs, in an unplanned scenario the cluster service will choose the appropriate destination server to move the VMs to. If a specific cluster node is found to not have sufficient resources such as CPU or memory to host the VMs, the cluster service will attempt to start the VM on another available node in the cluster.

 **NOTE:** Using the advanced features available in Windows Failover Clustering, rules may be set to restrict certain VMs to only run on certain nodes of a Hyper-V Host Cluster. For more information, refer to "Administering High Availability for Hyper-V" on page 45.

 **NOTE:** Because the state of the VM is not retained during an unplanned failover, the processors (and their features) on the originating and destination nodes do not have to match. However, Dell strongly recommends that all servers in a Hyper-V Host Cluster have identical hardware to ensure that it also supports planned migrations.

Comparison of Guest and Host Clustering

Table 2-1 highlights the key differences between Hyper-V Guest Clustering and Hyper- V Host Clustering.

Table 2-1. Comparison of Guest and Host Clustering

Category	Hyper-V Guest Clustering	Hyper-V Host Clustering
High Availability Scope	<ul style="list-style-type: none"> • The applications running on the guest OS are the clustered entities • HA is only for these applications and not for the VM itself 	<ul style="list-style-type: none"> • The virtual machines are the clustered entities • HA is for the virtual machines as a whole (including the guest OSes and applications within the VM)
Guest OS Requirements	<ul style="list-style-type: none"> • Requires that the guest OS is supported by Hyper-V • Requires that the guest OS has cluster software capability • Requires that this cluster software supports iSCSI storage 	<ul style="list-style-type: none"> • Solution is Guest OS independent because the cluster software is running in parent partition
Shared Storage Type	<ul style="list-style-type: none"> • iSCSI only • iSCSI initiator runs in the guest OS 	<ul style="list-style-type: none"> • Direct Attach SAS • iSCSI • Fiber Channel • SMB Share
Implementation	<ul style="list-style-type: none"> • Clustering has to be configured within every VM • May be more complex to manage if several VMs need to be made into HA VMs 	<ul style="list-style-type: none"> • Clustering has to be configured only in the parent partition on each host • Implementation is independent of the number of VMs

Table 2-1. Comparison of Guest and Host Clustering (continued)

Category	Hyper-V Guest Clustering	Hyper-V Host Clustering
Application support	Workloads within the VM must be cluster-aware	Workloads within VMs do not have to be cluster-aware
Maximum number of physical servers in the cluster	No direct dependency on physical servers. Limitations may apply based on the configuration.	Maximum of 16 physical servers in a single cluster configuration
Quick Migration	Not Available	Supported



NOTE: Based on the HA requirements, these models may be implemented in combination. However, careful planning and consideration is needed before implementation because such an implementation will be fairly complicated to configure and manage.

Designing for High Availability

Planning and designing is a very crucial step in ensuring a highly available environment that also provides the most isolation and best performance for the HA solution. This section provides guidance on a few of the key scenarios in failover cluster configuration.

Node and Disk Majority versus Node Majority Cluster Configuration

As mentioned earlier, Hyper-V is dependent on the Failover Clustering feature to provide HA capabilities. The Failover Clustering feature in Windows Server 2008 offers multiple configuration models for clustering up to 16 physical servers. The difference between the clustering models is in which of the cluster resources get to "vote" to decide whether or not the cluster is in a highly available state at any given time. The most common and recommended cluster configuration models are:

- Node and Disk Majority Cluster
- Node Majority Cluster



NOTE: A third type called Node and File Share may be suitable for certain scenarios. However, that is beyond the scope of this document. For information on all the cluster models, please refer to Microsoft website at www.microsoft.com/windowsserver2008/failover-clusters.msp.

Node Majority Cluster is a configuration in which only the servers that are part of the cluster can vote, and Node and Disk Majority is a configuration in which the servers and a single shared disk (SAN) resource get to vote. The choice of the cluster model is dependent on the specific requirements of the environment that the cluster is being deployed in.

Table 3-1. Comparison of Node Majority versus Node and Disk Majority

	Node Majority	Node and Disk Majority
Number of nodes	Recommended for clusters with an odd number of server nodes	Recommended for clusters with an even number of server nodes

Table 3-1. Comparison of Node Majority versus Node and Disk Majority (continued)

	Node Majority	Node and Disk Majority
Shared Storage	<ul style="list-style-type: none">• Shared storage is not required for configuring the Failover Cluster• HA VMs will require shared storage	<ul style="list-style-type: none">• Minimum 500MB NTFS shared disk as Witness Disk for configuring the Failover Cluster• Shared Storage is required for HA VMs
Requirement to Maintain High Availability Status	More than 50% of the nodes must be up for the cluster to be running	<ul style="list-style-type: none">• At least 50% of the nodes and witness disk must be available for the cluster to be running• If the witness disk is not available, more than 50% of the nodes must be available for the cluster to be running

Storage Configuration

The *Dell Storage Solutions Guide for Microsoft Hyper-V* provides detailed information on the various storage configuration options in Hyper-V as well as best practices for implementing them on Dell storage arrays. This section discusses the storage considerations in a Hyper-V Host Cluster environment.

Sharing a LUN for multiple HA VMs

A key decision point for Hyper-V Host Clusters is whether each HA VM will have its own LUN or if multiple HA VMs will share a single LUN. The main difference between these two configurations is in how they handle failovers.

The NTFS file system grants access at a LUN or volume level and not a file level. So when an HA VM is moved across physical servers, the ownership of all the storage LUNs that host that HA VM's files will also transfer across physical servers. This behavior impacts both planned and unplanned migrations.

- In planned migration (quick migration) scenarios, if there is a need to quick migrate an HA VM independent of other HA VMs, then that HA VM cannot share any of its LUNs with other HA VMs.

- In unplanned migration scenarios, this might be a factor depending on the resource requirements for all the HA VMs that share a LUN. If the cluster service is unable to successfully bring all the HA VMs (that share the single LUN) online on another node, it will retry on all other available cluster nodes. If none of the available cluster nodes can host all the HA VMs due to resource constraints, it is likely that none of the HA VMs that share the single LUN will be brought online

Table 3-2 shows important factors to consider when planning to share a LUN across multiple HA VMs:

Table 3-2. Determining if a LUN can be Shared

Single LUN can be shared by multiple HA VMs	Single LUN is required per HA VM
<ul style="list-style-type: none"> • If the HA VMs are only using VHD-based virtual hard disks (and not passthrough disks) 	<ul style="list-style-type: none"> • If the HA VMs are using passthrough disks (and not VHD-based virtual hard disks)
and	or
<ul style="list-style-type: none"> • If it is acceptable that all HA VMs will move across hosts simultaneously in planned failover scenarios (Quick Migration) 	<ul style="list-style-type: none"> • If there is requirement to be able to migrate HA VMs one at a time in planned failover scenarios (Quick Migration)



NOTE: Depending on the Disk I/O in the specific configuration, sharing a single LUN across multiple HA VMs may also impact I/O performance.

Storage Options for Hyper-V Host Clustering

The type of storage that is used in a Hyper-V Host Cluster will most impact the design of the Hyper-V Host Cluster solution. It will affect:

- The number of network adapters on each Hyper-V host
- The number of storage adapters on each Hyper-V host
- Quick Migration because VM state is saved to disk over the storage fabric

Table 3-3 summarizes the different types of storage that support Hyper-V Host Clustering and Hyper-V Guest Clustering. For a full list of supported storage hardware for Hyper-V, refer to the *Dell Storage Solutions Guide for Microsoft Hyper-V*.

Table 3-3. Storage Types for Hyper-V Clustering

Storage Type	Examples	Support for Hyper-V Host Clustering	Support for Hyper-V Guest Clustering
External Enclosures (JBODs)	PowerVault MD1000, PowerVault MD1120	No	No
SAS Arrays	PowerVault MD3000	Yes	No
iSCSI Arrays	PowerVault MD3000i, EqualLogic PS series, Dell EMC AX4-5	Yes	Yes
Fiber Channel Arrays	Dell EMC CX series	Yes	No

Network Configuration

The *Dell Networking Solutions Guide for Microsoft Hyper-V* provides detailed information on the various networking configuration options in Hyper-V as well as implementation best practices. This section discusses the networking considerations for a Hyper-V Host Cluster environment.

Cluster Network Types

In addition to storage resources, the cluster service also manages network resources on hosts that are part of the cluster. The cluster utilizes and manages these network resources based on their **cluster use** type as designated by the administrator when configuring the failover cluster. The **cluster use** type may be set by using the **Failover Cluster Management** console and selecting the specific adapters listed under **Network**. There are three **cluster use** types:

- **Internal Type:** The cluster network resource is configured as **Allow the cluster to use the network** and the option to **Allow clients to connect through the network** is NOT selected.
- **Enabled Type:** The cluster network resource is configured as **Allow the cluster to use the network** and the option to **Allow clients to connect through the network** option IS selected.
- **Disabled Type:** The cluster network resource is configured as **Do not allow the cluster to use this network**.



NOTE: These network types are defined by the Windows Failover Clustering feature in Windows Server 2008 and are not related to the Hyper-V network types that are defined in the *Dell Networking Solutions Guide for Microsoft Hyper-V*.

In a recommended configuration, every node of a typical Hyper-V Host Cluster should have dedicated network ports for the following:

Table 3-4. Dedicated Network Ports for Hyper-V Host Cluster Nodes

Network Types	Description, Configuration, and Best Practices
Network for Cluster Private	<ul style="list-style-type: none"> • One Gigabit (or Fast Ethernet) Ethernet network port for hosting the Cluster Private network. • This network is dedicated for use by the cluster service to check the health state of each cluster node. • Should be of cluster use type Internal.
Network for Cluster Public and for Parent Partition Management	<ul style="list-style-type: none"> • One Gigabit Ethernet network port for hosting the Cluster Public network and for network access to the Windows Server 2008 parent partition. • This network will be used to: <ul style="list-style-type: none"> – Manage the Hyper-V Host Cluster using the Failover Management console – Manage the Parent Partition and the Hyper-V Role • Should be of cluster use type Enabled.
Network for Virtual Machines	<ul style="list-style-type: none"> • One or more Gigabit Ethernet network ports to host virtual machine networks. • Should be of cluster use type Disabled to ensure isolation of the virtual machine network traffic.
Network for Storage	<ul style="list-style-type: none"> • Two or more Gigabit Ethernet network ports to host iSCSI storage network (for iSCSI-based configurations). • Should be of cluster use type Disabled to ensure isolation of the storage network traffic.



NOTE: Even though Failover Cluster in Windows Server 2008 does not require a dedicated cluster private network because it can use the public network to monitor cluster heartbeat, Dell recommends a dedicated cluster private to ensure the reliability of the Hyper-V Host Cluster configuration.

Best Practice Recommendations

- Static IP addresses should be used for networks — especially the Cluster Private, Cluster Public, Parent Partition, and any iSCSI Storage networks.
- Ensure that the Cluster Private network is on a separate subnet than all other networks on the Hyper-V hosts. It is recommended that each network type in Table 3-4 is on a separate subnet.
- Ensure that the network configuration is completed (including assigning the appropriate subnets and static IP addresses) on every Hyper-V host before the Hyper-V Host Cluster is configured.



NOTE: This is critical because in Windows Server 2008, the Failover Cluster service automatically identifies network resources during configuration and assigns them as cluster network resources. It does this based on the subnet configurations and automatically selects at least one network adapter from each subnet. Network resources cannot be manually added or removed as a cluster resource using the **Failover Cluster Management** console

- After the Failover Cluster is configured, ensure that cluster network resources have been configured as recommended in Table 3-4.
- When using dual-port or quad-port network adapters, host the same network types on different physical network adapters to ensure that there are redundant network paths available to accommodate a single network adapter (NIC) HW failure. For example, in a configuration with two iSCSI storage networks, instead of using two ports of the same network adapter for hosting the iSCSI network, use network ports on two separate network adapters. Follow similar best practices for Cluster Private/Public, and Virtual Machine Networks.

Recommended Configurations

As mentioned in the *Dell Solutions Overview Guide for Microsoft Hyper-V*, Dell supports Hyper-V on a very wide array of server, networking, and storage hardware. Even though the Failover Clustering feature in Windows Server 2008 does not require that the hardware in a cluster configuration be identical, Dell strongly recommends having identical hardware in any single Hyper-V Host Cluster configuration. This will ensure that there are no compatibility issues (such as processor dependencies to support Quick Migration) that could cause disruption in a production environment.

Dell offers a wide array of servers that have been specifically built with the goal of being able to host virtualized workloads. These servers offer support for large amounts of physical memory, large numbers of I/O slots for both network and storage, and a minimal power footprint. For a full list of servers that Dell recommends for virtualization, go to www.dell.com/servers.

Design for iSCSI-Based Configurations

Table 3-5 summarizes the recommended network requirements for an iSCSI-based configuration for a Hyper-V Host Cluster:

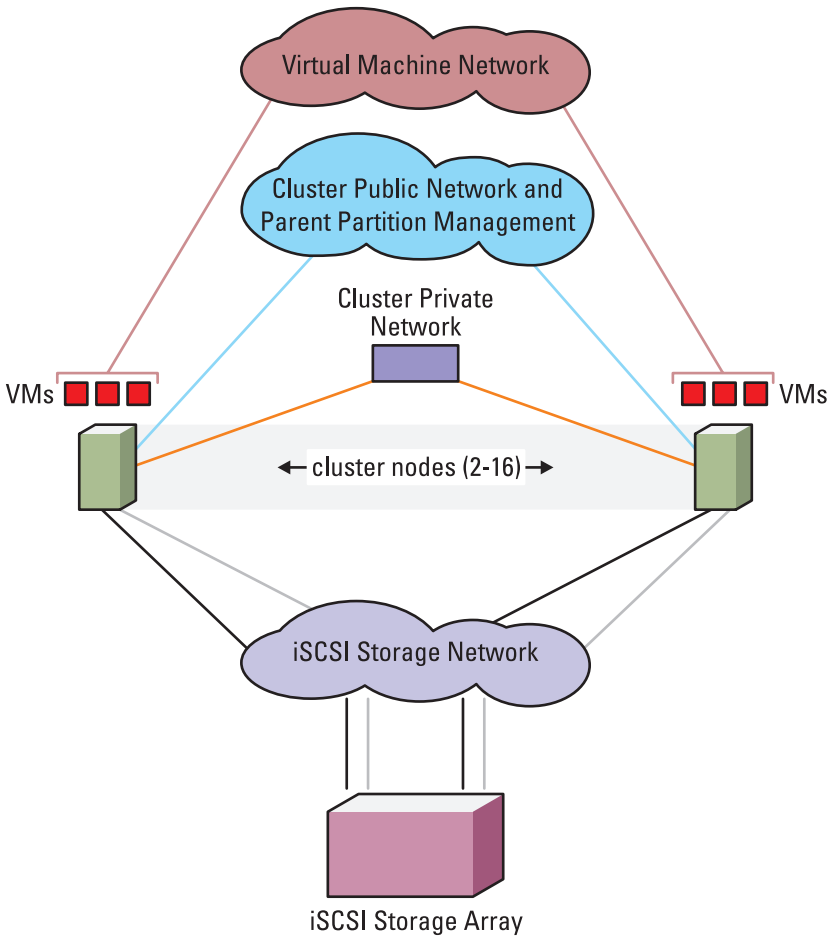
Table 3-5. iSCSI-Based Configuration Network Requirements

	Parent Partition Management & Cluster Public Network	Cluster Private Network	iSCSI Storage Network	Virtual Machine Network
Minimum Recommended Ports	1 Port	1 Port	2 Ports	2 Ports*

* Scale the number of network ports based on the number of VMs and workloads. Distribute VMs across the associated virtual networks based on bandwidth requirements.

Figure 3-1 provides a sample configuration of the infrastructure and highlights the need for isolating the traffic for management, storage, and workloads running within the virtual machines.

Figure 3-1. iSCSI-Based Hyper-V Host Cluster Configuration



Design for Fiber Channel Configurations

Table 3-6 summarizes the recommended network requirements for a Fiber Channel based configuration for a Hyper-V Host Cluster:


Table 3-6. Fiber Channel-Based Configuration Network Requirements

	Network Connections			Fiber Channel Connections
	Parent Partition Management and Cluster Public Network	Cluster Private Network	Virtual Machine Network	Storage I/O
Minimum Recommended Ports	1 Port	1 Port	2 Ports*	2 Ports

* Scale the number of network ports based on the number of VMs and workloads. Distribute VMs across the associated virtual networks based on bandwidth requirements.

Figure 3-2 provides a sample configuration of the infrastructure and highlights the need for isolating the traffic for management, storage, and workloads running within the virtual machines.

Figure 3-2. Fiber Channel-Based Hyper-V Host Cluster Configuration



Design for Direct-Attach SAS Configurations

Table 3-7 summarizes the recommended network requirements for a Direct-Attach SAS-based configuration for a Hyper-V Host Cluster. It is important to note that this is only guidance, and actual deployment configuration will vary based on the existing infrastructure.

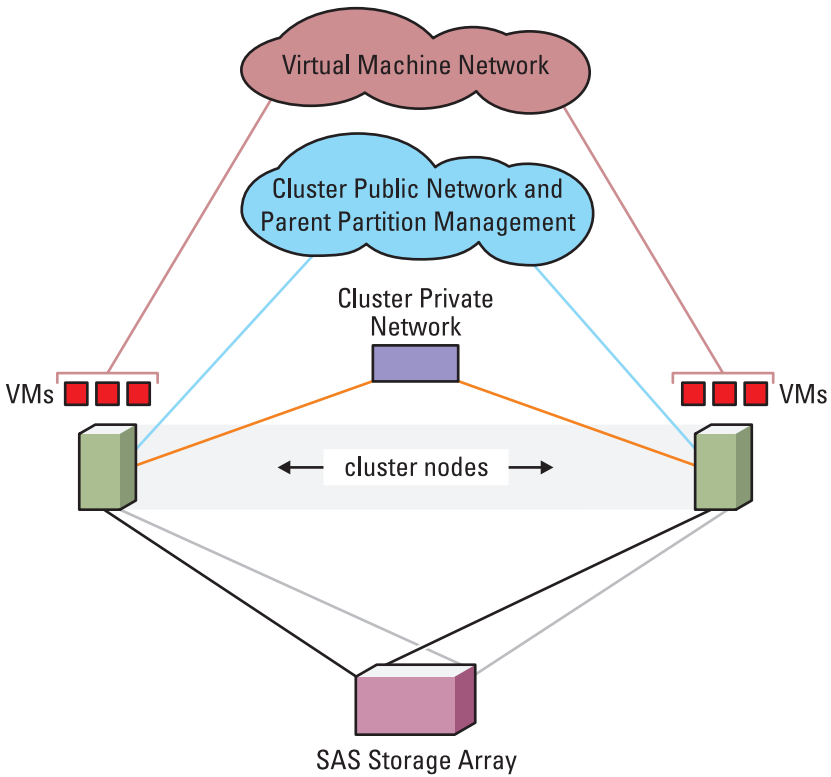
Table 3-7. Direct-Attach SAS-Based Configuration Network Requirements

	Network I/O Connections			SAS Connections
	Parent Partition Management and Cluster Public Network	Cluster Private Network	Virtual Machine Network	Storage I/O
Minimum	1 Port	1 Port	2 Ports*	2 Ports
Recommended				
Ports				

* Scale the number of network ports based on the number of VMs and workloads. Distribute VMs across the associated virtual networks based on bandwidth requirements.

Figure 3-3 provides a sample configuration of the infrastructure and highlights the need for isolating the traffic for management, storage, and workloads running within the virtual machines.

Figure 3-3. SAS-Based Hyper-V Host Cluster Configuration



NOTE: Even though it is possible to configure a 4-node Hyper-V Host Cluster with a Dell MD3000 SAS Storage Array, Dell recommends only a 2-node Hyper-V Host Cluster with the MD3000 SAS Array. This is to ensure that the configuration is more reliable. In a 4-node configuration, there is only one path from any one node to the storage array and that does not provide multi-path resiliency for the solution.

Integration between Hyper-V and Windows Failover Clustering

As mentioned earlier, Hyper-V HA is implemented by leveraging the Windows Failover Clustering feature in Windows Server 2008. A high level knowledge of Failover Clustering is essential for implementing and managing a Hyper-V Host Cluster environment. Refer to

<http://www.microsoft.com/windowsserver2008/en/us/clustering-overview.aspx> for an overview of the Failover Clustering feature.

In order for Failover Clustering to manage an application as a highly available entity, it has to manage all the dependencies for that application as well. Failover Clustering achieves this by defining a logical entity called a Resource Group and placing the application and its dependent resources within that Resource Group. When a failover occurs, Failover Clustering moves all the resources defined within a Resource Group (as a whole) across cluster nodes. In the **Failover Management** console for Failover Clustering, the Resource Group and its contents are listed under the **Services & Applications** section.


Failover Clustering extends this same concept to Hyper-V Host Clustering as well. When a VM is made an HA VM, Failover Clustering creates a new Resource Group for this VM. The logical Resource Group for an HA VM will, at a minimum, include the following resources:


- The **Virtual Machine** resource defines the state of the HA VM.
- The **Virtual Machine Configuration** resource maintains the HA VM's configuration. This resource has to be online for the VM to be available on **Hyper-V Manager**.
- The **Disk Drives** that host the XML, BIN, VSV, and VHD files (or passthrough disks).

The Resource Group may include more storage disks in more complex configurations such as when the VM has multiple disks (including passthrough disks).

With its integration with Hyper-V, the Failover Clustering feature can automatically identify all the dependencies for a VM and groups all the identified resources into a single Resource Group. This activity is transparent


to the administrator. At a high level, Failover Clustering goes through the following process when an administrator chooses to make a VM into a HA VM:

- 1 Parses the VM's configuration file to identify the dependent resources
 - 2 Checks whether all the dependent disks (hosting VM files, VHD files, or Passthrough disks) have been configured as shared storage for all nodes of the cluster
 - 3 Checks whether all the dependent resources (such as any files shares) are accessible by all nodes of the cluster
-  **NOTE:** Even though it is a requirement that shared storage disks are cluster resources, the process does not explicitly check that they are.
- 4 Adds the Virtual Machine, Virtual Machine Configuration, and Storage Resources to a newly defined Resource Group
 - 5 Upon completion, generates a report that provides details of the HA configuration and, most importantly, calls out any warnings for the HA VM configuration

 **NOTE:** It is strongly recommended that all the warnings be addressed before proceeding with the HA VM configuration, especially in a production environment.

Microsoft KB951308 QFE for Hyper-V Host Clusters

KB951308 is an update that addresses several critical issues for Hyper-V Host Clustering and implements better integration between the Hyper-V Role and the Failover Clustering feature in Windows Server 2008. Visit the Microsoft website at <http://support.microsoft.com/kb/951308> to download this QFE and to get important information on the functionality enabled by the QFE. Make sure to download the x64 version of the QFE.

 **NOTE:** Dell strongly recommends that this QFE be applied on all Hyper-V cluster nodes as well as systems running the Failover Clustering MMC. This solutions guide assumes that QFE 951308 has been applied on all the cluster nodes as well as systems running the failover cluster management console.

Implementing High Availability in Hyper-V

This section covers the procedures for implementing High Availability in Microsoft Hyper-V. This document only provides high level guidance. For detailed instructions, please refer to *Failover Clustering in Windows Server 2008* at www.microsoft.com/windowsserver2008/failover-clusters.mspx.




Implementing Hyper-V Host Clusters

For any successful Hyper-V Host Cluster deployment, careful consideration must be given to the network and storage configurations to ensure that there are no single points of failure and to ensure that there is sufficient bandwidth for both network and storage traffic. For more information, see "Designing for High Availability" on page 21.



This section discusses best practices for configuring a Hyper-V Host Cluster and provides high-level steps to setup a Hyper-V Host cluster. For step-by-step instructions on configuring a Hyper-V Host Cluster, refer Microsoft documentation *Hyper-V Step-by-Step Guide: Hyper-V and Failover Clustering* at <http://technet.microsoft.com>.

Preparing the Systems

- 1 Install Windows Server 2008 x64 Enterprise or DataCenter Editions on your servers.
- 2 Make sure that the Hyper-V RTM upgrade patch (KB950050 from the Microsoft website) has been applied on all servers.
- 3 Make sure that the 951308 QFE (KB951308 from the Microsoft website) has been applied on all servers including systems that will run the **Failover Cluster Management** console (cluadmin.msc). Make sure to download the x64 version of the QFE.

- 4 Make sure that all the devices on the server have the latest supported drivers from <http://support.dell.com>. Download the device drivers for Windows Server 2008 x64.
 -  **NOTE:** Dell does not support device drivers that are downloaded from 3rd party vendors. Only use drivers from <http://support.dell.com>.
- 5 Configure the network connections and storage connections as planned.
 -  **NOTE:** Make sure that the appropriate MPIO software for the storage array has been installed on the nodes.
 -  **NOTE:** For Node and Disk Majority configuration of the Failover cluster, make sure that the shared LUN (minimum 500 MB) that will be used as the witness disk is configured as shared storage for all servers.

Configure the Failover Clustering Feature and Hyper-V Role

- 1 Install the Failover Clustering feature on all the Windows Server 2008 servers that will be part of the Hyper-V Host Cluster. The feature can be enabled using the **Add Feature** wizard in **Server Manager**.
- 2 The Failover Clustering feature in Windows Server 2008 provides a new capability that allows administrators to run a validation check on the planned configuration to identify potential issues before actually creating the failover cluster. Open the **Failover Cluster** console on any of the systems, run the **Validate a Configuration** wizard, and make sure that the configuration with all the nodes meets all the requirements for a Windows Failover Cluster.
 -  **NOTE:** It is strongly recommended that all errors and warnings reported by the validation wizard are addressed before proceeding with configuring the failover cluster.
- 3 After all the requirements have been verified to be met, use the **Failover Cluster Management** console to create the failover cluster.
- 4 Enable the **Hyper-V Role** on each of the nodes of the failover cluster.
 -  **NOTE:** For information on enabling the Hyper-V Role, refer to the *Dell Solutions Overview Guide for Microsoft Hyper-V*.
- 5 Configure the Virtual Networks as planned on each Hyper-V server. Make sure that Virtual Networks are created with the exact same names on all the servers in the Hyper-V Host Cluster. For example, if any one cluster node has a virtual network called `TestNetwork`, all the Hyper-V hosts in

the cluster should have a Virtual Network with the name `TestNetwork`. If not, when a HA VM fails over to a system that does not have a Virtual Network with that name, the VM will not have network connectivity. In such a scenario, the network connectivity will have to be enabled by manually configuring the virtual network for the HA VM.

Creating Virtual Machines that will be HA VMs

This section discusses the steps to configure a new virtual machine that will be made into an HA VM. For configuring existing non-HA VMs as HA VMs, please refer to "Converting an Existing Non-HA VM to a HA VM" on page 49. Choosing the storage location that will host the virtual machine's files is the most important aspect of configuring a HA VM. As mentioned earlier, Hyper-V requires that all the files and disks (VHD or Passthrough) belonging to a HA VM are on shared storage (SAN) and are accessible to all nodes of the Hyper-V Host Cluster. For creating an HA VM, the following consoles will be required:

- **Failover Cluster Management** console (`cluadmin.msc`)
- **Hyper-V Manager** console (`virtmgmt.msc`)
- **Disk Management** console (`diskmgmt.msc`)
- Storage array software console such as:
 - Dell Modular Disk Storage Manager™ for PowerVault MD3000 and MD3000i arrays
 - Dell EqualLogic™ Group Manager for Dell EqualLogic arrays
 - EMC Navisphere Manager for Dell|EMC arrays

It is recommended that this process be done one VM at a time to ensure that storage disk allocations to VMs are correct. Storage configuration for a VM involves the following:

- 1 Use the SAN storage array software console to create the logical disks (LUNs) that are required for the VM. Make sure all the disks on the SAN are configured to be available to all the nodes of the cluster.




NOTE: For guidance on sizing storage for virtual machines, refer to the *Dell Storage Solutions Guide for Microsoft Hyper-V*.


- 2 Use the **Disk Management** console on each of the server nodes to verify that the newly provisioned storage LUN is visible to all nodes.

 **NOTE:** The LUN may be in offline state at this time on all the hosts.


- If the LUN will be used as a passthrough disk for the VM, make sure to bring the LUN **Online** on any one cluster node, **Initialize** it, and then take it **Offline**.
- If the LUN will host the VM's VHD file, bring the disk **Online** on any one cluster node, **Initialize** it, and **Format** it as NTFS.

 **NOTE:** In scenarios when the Hyper-V Host Cluster is hosting a relatively large number of VHD-based VMs, the cluster configuration may run out of drive letters in the parent partition. Hyper-V allows the use of Mount Points to host storage LUNs to work around the drive letter limitation. Refer to "Using Mount Points in Hyper-V" on page 55 for more information.

- 3 Use the **Failover Cluster Management** console to add the storage LUNs as cluster resources. Confirm that the newly provisioned storage disks are listed as **Storage**→**Available Storage** disks and are not allocated to any specific Resource Groups.

 **NOTE:** If the storage disks are not made cluster resources at this stage, the attempt to convert a VM to an HA VM may result in an improper configuration (this is expected behavior after QFE 951308 is applied).

- 4 When a storage disk is added as a cluster resource, the cluster will choose to bring the storage disk online on any one of the cluster nodes.
- 5 Use the **Failover Cluster Management** console to identify the host that owns the newly provisioned storage LUN.
- 6 Use the **Hyper-V Manager** console on that specific host to launch the **New Virtual Machine** wizard.

 **NOTE:** If the VM creation has to be done on a cluster node other than the cluster node that owns the newly provisioned storage LUN, then the storage LUN has to be moved to the node that you are attempting to create the VM on. Cluster disks in the **Available Storage** group can only be moved across cluster nodes using the CLI. For example, the CLI command `cluster group "Available Storage" /move` can be used to move storage disks across nodes in a two-node cluster.

- 7 In the **Specify Name and Location** screen of the **New Virtual Machine** wizard, select **Store the virtual machine in a different location** and specify the path to the shared storage disk that will host the virtual machine files (XML/BIN/VSV).



NOTE: If the VM's boot disk will be a passthrough disk, the VM's files (XML/BIN/VSV) can either be on another shared storage LUN or a file share. See "Using SMB/CIFS Shares in Hyper-V" on page 53 for more information on file shares.

- 8 At the **Connect Virtual Hard Disk** screen, choose the type of boot disk for the VM (**VHD** or **Passthrough**). For **VHD**, specify the path to the shared storage LUN. For **Passthrough** disks, select **Attach a virtual hard disk later** and attach the disk later by updating the hardware configuration of the VM.
- 9 Follow the prompts of the wizard to complete the VM configuration.
- 10 The VM is now ready to be configured as an HA VM. The administrator may choose to install the guest operating system at this time.



NOTE: Before a VM can be made an HA VM, the VM's DVD Drive should be detached from physical optical drives or image files. If Integration Services were installed in the guest operating system, make sure the Integration Services CD image is detached from the VM's DVD drive. If there is a specific requirement that a VM's DVD Drive has to be in captured state, then make sure the DVD Drive is attached to an image file via a UNC path that is accessible to all nodes of the cluster.

Enabling VMs as HA VMs

Virtual machines created on a host that is a part of a Hyper-V Host Cluster do not automatically become HA VMs. The administrator has to choose to make VMs into HA VMs.

- 1 Make sure the VM that will be made an HA VM is either turned off or shut down.
- 2 Launch the **Failover Cluster Management** console for the system that is currently running the VM.
- 3 On the console, click **Configure a New Service or Application** to launch the **High Availability Wizard**.

- 4 When prompted by the wizard, choose **Virtual Machine** as the resource type and the wizard will automatically list all the VMs that meets the basic requirements for an HA VM.



NOTE: If a specific VM is not listed here, make sure that VM has met all the pre-requisites for HA VMs.

- 5 Select the appropriate VM and proceed with the configuration.
- 6 Upon completion, the **High Availability** wizard also generates a report of the HA VM configuration that needs to be reviewed to make sure there were no issues.
- 7 The VM is now an HA VM and is ready to be turned on from either the **Hyper-V Manager** console or the **Failover Management** console.
- 8 At the end of this configuration, the **Failover Cluster Management** console will list the newly created Resource Group, `Virtual Machine (n)`, for the HA VM. This `Virtual Machine (n)` Resource Group will list all the dependencies for the VM.



NOTE: It is recommended to run the dependency report for the Resource Group to confirm that the cluster service has made the right dependency mappings for the HA VM. Click on the "Show Dependency Report" available in the **Failover Clustering** console to generate this dependency report.

- 9 It is recommended the name of the resource group be changed from `Virtual Machine (n)` to a name that more closely reflects the VM hosted within that group.

Migrating HA VMs Across Hosts in the Cluster

With the VMs configured as HA VMs, they are now capable of handling both planned and unplanned downtime scenarios. For unplanned downtime, the Failover Cluster service will take care of moving the HA VMs on the other cluster nodes. For planned downtime scenarios, administrators can migrate HA VMs to the other cluster nodes.

This activity is initiated from the **Failover Cluster Management** console.

- 1 Open the **Windows Failover Cluster Management** console and select the appropriate **Resource Group** hosting the HA VM that needs to be migrated to another node.
- 2 Select **Move virtual machine(s) to another node**. The console will list the available target nodes in the Hyper-V Host Cluster.

- 3 Select a target node and the HA VM will be migrated to and started on the target node automatically.

Hyper-V Host Clusters on Server Core

Windows Server 2008 offers a new installation option called Server Core that has a smaller footprint installation on the operating system. The key benefits of choosing a Server Core installation option include:

- Reducing the amount of future patching due to its reduced footprint
- Reducing attack surface because highly targeted roles such as Internet Explorer and Explorer Shell are not available in a Core environment

Server Core does not present a local GUI interface. This may introduce some new challenges because managing the system locally can only be done through a Command Line Interface (CLI). However, once a server running Server Core has been set up for remote administration, all the Roles and Features installed on that server (including Hyper-V and Failover Clustering) can be managed remotely using their respective GUI-based MMC consoles.

Setting up the Server Core System for Hyper-V and Failover Clustering

After the Server Core OS installation is complete, refer to the following documentation for guidance on the initial configuration tasks as well as enabling the system for complete remote management:

- *Dell Solutions Overview Guide for Microsoft Hyper-V* available at www.dell.com/hyper-v
- *Server Core Installation Option of Windows Server 2008 Step-By-Step Guide* available on the Microsoft website
- *Server Core Remote Management* at <http://technet.microsoft.com>

To enable the Hyper-V Role and Failover Clustering features on all the Server Core nodes:

- 1 Use `start /w ocsetup Microsoft-Hyper-V` to install the Hyper-V Role.
- 2 Use `start /w ocsetup FailoverCluster-Core` to install the Failover Clustering feature.

- 3 Install the Hyper-V RTM update (KB950050 from the Microsoft website) for Windows Server 2008 x64 Editions. Note that the update is the same for Full and Server Core installations.
- 4 Install the Failover Clustering QFE 951308 (KB951308 from the Microsoft website) on Server Core. The QFE bits are the same for Full and Server Core installations. Make sure to download the x64 version of the QFE.

Refer to the following documents for configuring networking and shared storage on the Server Core systems:

- *Dell Networking Solutions Guide for Microsoft Hyper-V* available at www.dell.com/hyper-v
- *Dell Storage Solutions Guide for Microsoft Hyper-V* available at www.dell.com/hyper-v

Setting up Remote Management to Server Core

At a minimum, the following GUI consoles are required to remotely manage a Hyper-V Host Cluster:

- Hyper-V Manager (virtmgmt.msc)
- Failover Clustering Management (cluadmin.msc)

Both these consoles can be installed either on a Full install of Windows Server 2008 (32-bit or x64) or Vista (32-bit or x64) to remotely manage Hyper-V Host Clustering on Window Server 2008 Server Core.

Remote Management from Windows Server 2008 Full install

- 1 Install the KB950050 Hyper-V RTM update for Windows Server 2008. This update has to be applied even on systems that are only running the Hyper-V Manager console. Separate packages are available for 32-bit and x64 systems.
- 2 Install the Failover Clustering QFE 951308 (KB951308 from the Microsoft website).
- 3 On the remote Windows Server 2008 system (32-bit or x64), open **Server Manager** (servermanager.msc), and click **Add Features Wizard**.
- 4 Under **Select Features**, choose the following:
 - **Hyper-V Tools** under **Role Administration Tools**.
 - **Failover Clustering Tools** under **Feature Administration Tools**.

- 5 Proceed with the installation of the selected features.
- 6 After installation is complete, these consoles are available under **Administrative Tools**.

Remote management from Windows Vista

- 1 Install the Remote Server Administration Tools (RSAT) for Windows Vista available for download from the Microsoft website. Separate packages are available for the x86 and x64 versions of Vista. RSAT is a collection of tools that are meant to allow administrators to remotely manage the Windows server environments including Windows Server 2008.
- 2 To install the **Hyper-V Manager MMC** console on Vista, download the appropriate 32-bit or x64 bits for KB952627.
- 3 Install the Failover Clustering **QFE 951308** (KB951308 from the Microsoft website).
- 4 Launch an MMC console (**Start→Run** and type **MMC**) and add the **Failover Cluster Management** and **Hyper-V Manager** snap-ins to the MMC console.

Administering High Availability for Hyper-V

This section provides guidance on the common administrative tasks to manage HA VMs in a Hyper-V Host Cluster environment.

Changing the configuration of HA VMs



NOTE: The QFE in Microsoft KB951308 is required for the functionality described in this section. For more information, see "Microsoft KB951308 QFE for Hyper-V Host Clusters" on page 34.

Changes to a virtual machine configuration may or may not involve changing the hardware profile of the VM. The steps required to successfully handle a configuration change for an HA VM differs between whether the VM's hardware profile is changed or not.

For Configuring Changes that Impact the VM's HW Profile

These are the changes that require the VM be in an OFF state or SHUTDOWN state. Examples of this type of change include changing the number of virtual processors, changing virtual memory, adding/removing network and storage controllers, etc. Here are the high-level steps to make this type of configuration change for an HA VM:

- 1 To successfully shut down or turn off an HA VM, use the VM control actions available in the Failover Clustering MMC. Do not choose these actions from the Hyper-V Manager console because that will cause the cluster service to automatically restart the HA VM to maintain its HA status.
- 2 Once the HA VM is shutdown, use the Hyper-V Manager Console to make the changes.
- 3 After the changes are complete, use the Failover Cluster console to start the HA VM.

For Configuring Changes that DO NOT Impact the VM's HW Profile

These are the changes that can be done when the VM is running. Examples of this type of change include changing the virtual network that a virtual network adapter is connected to, changing the media type that is attached to a VM's DVD drive, etc.

The Failover Cluster Management console provides an option on the right side **Actions** panel to **Refresh virtual machine configuration** for HA VMs. This option is meant to make sure that the cluster service has the most up-to-date configuration information of the HA VM. Activity such as **Quick Migration** of HA VMs may fail if the cluster service does not have the latest configuration information of the HA VM. The **Action** option should be chosen whenever any changes are made that **DO NOT** impact the HA VM's HW profile.

Here are the high-level steps to make this type of configuration change for an HA VM

- 1 Use the Hyper-V Manager console to update the HA VM's configuration information.
- 2 Use the Failover Cluster Management console to "Refresh virtual machine configuration" for that specific VM.



NOTE: The virtual machine will be taken offline for this action.

- 3 Validate the report generated by the wizard to confirm that the changes were successful.

Adding Passthrough Disks to HA VMs



NOTE: The QFE in Microsoft KB951308 is required for the functionality described in this section. For more information, see "Microsoft KB951308 QFE for Hyper-V Host Clusters" on page 34.

There are two prerequisites for adding additional passthrough disks:

- The storage disk needs to already be a cluster resource as an **Storage**→**Available Storage** in the **Failover Cluster** console.
- The storage disk resource has to be owned by the same cluster node on which the HA VM is currently running.

Ownership of HA resources can be verified using the **Failover Cluster** console. Check the Resource Groups listed under **Services and Applications** to verify ownership of HA VMs and check the **Storage** section to verify ownership of disk resources. If the physical disk resource listed under **Available Storage** and the virtual machine resource are owned by separate nodes, either of the resources needs to be moved so that they are both on the same node. If migrating the VM to the other node is not a viable option, then the storage resource may be moved. The latter is achieved via CLI. For example, the CLI command `cluster group "Available Storage" /move` can be used to move storage disks across nodes in a two-node cluster. Once the resources are owned by the same node, make sure the HA VM is in **Off** or **Shutdown** state and then proceed with adding the passthrough disk to the VM by updating the VM's hardware configuration.



NOTE: In order to enable the above functionality, Hyper-V allows passthrough disks that are already **Available Storage** cluster resources to be added an HA VM. However, currently Hyper-V does not check if a passthrough disk is already an allocated cluster resource before allowing it be added to an HA VM. However, currently Hyper-V does not explicitly check if a passthrough disk is an **Available Storage** cluster resource or whether it is an allocated cluster resource before allowing it be added to an HA VM. Make sure that the correct passthrough disk is being attached.

Sharing LUNs Between HA VMs

With the 951308 QFE applied, Windows Failover Cluster allows multiple HA VMs to be hosted within the same Resource Group. This is a critical requirement for scenarios such as when multiple HA VMs have VHD files on the same storage LUN. However, careful consideration needs to be given before multiple HA VMs are hosted within a single resource group because the resource groups failover as a single entity. All of the VMs hosted on a single resource group will failover simultaneously.

Multiple VMs can share a storage LUN only if the VMs are using VHD based rather than passthrough disks. Use the following guidelines to add a new VM to an existing HA VM group:

- 1 If a Resource Group has not been created already, use the procedures from "Implementing High Availability in Hyper-V" on page 35 to create an HA VM, for example HA-VM1.

- 2 Create a new VM, for example HA-VM2, and place its VHD file on the same storage LUN that was used by HA-VM1.
- 3 Use the **Windows Failover Cluster** console to make HA-VM2 also an HA VM.
- 4 The failover cluster service will automatically add HA-VM2 into the existing Resource Group.

HA VM Resource Group Names

By default, the cluster resource groups for HA VMs are created with generic names such as "Virtual Machine (1)" and does not inherit the names of the VMs. This is expected behavior after the 951308 QFE is applied and was designed to avoid naming conflicts in resource groups as well as to accommodate the scenario of a single resource group owning multiple HA VMs. However, administrators can manually change these Failover Cluster Resource Group names using the **Failover Clustering MMC**.

Cluster Resource Group for HA VMs Reported as "Partially Online"

When a Failover Cluster Resource Group contains more than one virtual machine and any of the virtual machines are not in the **Running** state, the status of the Resource Group will be displayed as **Partially Online**. Even a Resource Group with a single HA VM may report the status of **Partially Online** when any of the dependent resources of that HA VM (such as a passthrough disk that is not hosting the boot volume of the HA VM) is in an **Offline** state.

Configuring Rules for Associating Hosts to HA VMs

Failover Clustering in Windows Server 2008 allows the administrator to set rules on how cluster resources are managed among the nodes of the cluster. With Hyper-V HA VMs being managed by the Failover Cluster as a cluster resource, rules can be set for resource groups that host HA VMs as well.



NOTE: Dell strongly recommends not changing the default settings except in special circumstances.

Setting Host Preferences for HA VMs

This capability in Failover Clustering allows the administrator to specify the order of preference for cluster nodes that will host an HA VM. To set host preferences at a resource group level:

- 1 Open the **Failover Clustering** console, select **Services and Applications**, and right-click the Resource Group that hosts the HA VM.
- 2 Select the **General** tab and set the order of preferences in the **Preferred Owners** list.



NOTE: The **Failover** tab allows the administrator to specify the failback behavior, which is based on the **Preferred Owners**. It is recommended that failback be set to the default setting of **Prevent Failback** in a production environment.

Blocking HA VMs from Running on Certain Hosts

This capability in Failover Clustering allows the administrator to block certain hosts from hosting a specific HA VM. To block HA VMs at a resource group level:


- 1 Open the **Failover Clustering** console, select **Services and Applications**, and select the Resource Group that hosts the HA VM.
- 2 Right-click the **Virtual Machine** resource, select the **Advanced Policies** tab, and de-select the cluster nodes that should not host this specific virtual machine.
- 3 Repeat step 2 for all the other resources (**Virtual Machine Configuration** and **Disk Drives**) in this resource group.

Converting an Existing Non-HA VM to a HA VM

The key requirement for a VM to be converted to a HA VM is that all the VM's files (including VHDs) are on shared storage. "Implementing Hyper-V Host Clusters" on page 35 walks through the high-level steps for configuring HA VMs if the VM's files are already on shared storage and the disks can be made cluster resources.


The following is high level guidance for converting a non-HA VM that has its files on a disk that cannot be made a Windows Server 2008 Failover Cluster resource such a local disk on a server or a disk on an array that does not support Windows Server 2008 Failover Clustering. In this scenario, the VM's

files have to be moved from the current location to storage LUN(s) that can be made cluster resources. The recommended mechanism to move VM files is to use the **Export/Import** functionality in Hyper-V.

- 1 Provision a new storage LUN, bring the disk **Online** on any one cluster node, **Initialize** it, and **Format** it as **NTFS**.
- 2 Use the **Failover Cluster Management** console to make this storage LUN a cluster resource. Make sure that this LUN is listed under **Storage**→**Available Storage**.
- 3 From the **Hyper-V Manager Console**, select the non-HA VM, and **Export** the VM files to the newly provisioned storage LUN.
 **NOTE:** If the storage LUN cluster resource is not owned by this node, the storage LUN will need to be moved to the node that is currently running the VM before exporting the VM files.
- 4 After the **Export** is complete, delete the existing non-HA VM using the **Hyper-V Manager** console.
- 5 Use the **Import** functionality to import the VM files from the shared storage LUN to the same node.
- 6 Use the guidance provided under "Enabling VMs as HA VMs" on page 39 to complete configuration of the HA VM.

Converting an HA VM to a Non-HA VM

To convert an HA VM to a non-HA VM requires that the VM files be exported before the files are removed as cluster resources.

- 1 Use the **Failover Cluster** console to **Shutdown** or **Turn Off** the VM.
 **NOTE:** Do not initiate these actions from the **Hyper-V Manager** console because the cluster service will restart the VM to maintain the VM's HA state
- 2 From the **Hyper-V Manager** console, **Export** the virtual machine to an alternate location (such as a local hard disk). Once the export is complete, delete the existing VM from **Hyper-V Manager**.
- 3 Use the **Failover Clustering** console to delete the VM resource group that is hosting the HA VM. The VHD files have to be deleted manually from the shared storage.

- 4 From the **Hyper-V Manager** console, choose to **Import** the virtual machine from the location it was exported to in Step 2.



NOTE: If the resource group of an HA VM is deleted prior to exporting the VM files as recommended above, **Hyper-V Manager** will de-register the VM and the VM will not be listed on the console even though the files are still available on the disk. There is no mechanism to re-register the VM. In such a scenario, the only recovery option is create a new virtual machine from the **Hyper-V Manager** console and attach the VM to the existing VHD files or passthrough disks.

Using SMB/CIFS Shares in Hyper-V

Hyper-V allows the use of SMB/CIFS shares (file shares) to host VM resources such as ISO files or VHDs. Use of SMB/CIFS shares is not recommended due to the performance limitations but there may be certain scenarios that might benefit from their use. Consider the following when implementing SMB shares in a HA environment:

- Best practice recommendation when using an SMB share in Hyper-V HA environments is to make sure that the file share is also a clustered resource so that it does not become a single point of failure in the highly available environment.
- It is strongly recommended that the file share be managed by a cluster that is separate from the Hyper-V Host Cluster. If the file server resources are also hosted by the same cluster, there may be issues during failover scenarios due to the file server resource not being online when the dependent HA VM resources come online. This is especially true when the same cluster node ends up owning the file server resource and the dependent HA VM resource. For more information on configuring a file server cluster, download the Microsoft document *Step-by-Step Guide for Configuring a Two-Node File Server Failover Cluster in Windows Server 2008* from microsoft.com.
- When a VM that uses file shares is made a HA VM, the **High Availability Wizard** may generate a warning that the file shares are not clustered resources even when they are. This is known behavior and this warning may be ignored.
- When accessing SMB shares from the Hyper-V hosts via UNC path, best practice is to use the DNS name of the server hosting the share rather than the IP address of the server hosting the share. For example, when mapping to a shared folder **Share\$** on server **FileServer.Domain.com**, make sure that it is mapped using **\\FileServer.Domain.com\Share\$**.




NOTE: Hyper-V does not support access file share via a mapped drive.

- Grant the appropriate permissions for the domain "Machine" account of all the Hyper-V hosts that are part of the Hyper-V Host Cluster:
 - The domain machine accounts (such as DOMAIN\COMPUTERNAME\$) of all Hyper-V cluster nodes are granted **Full Control** in the **Share** permissions of the shared folder
 - The domain machine accounts of all Hyper-V cluster nodes are granted **Full Control** in the **Security** settings for the shared folder and all its files
- In certain scenarios (such as remotely managing a Hyper-V server), the use of SMB shares will require the use of constrained delegation. Refer to *Step-by-Step Guide for Configuring a Two-Node File Server Failover Cluster in Windows Server 2008* from microsoft.com.
- Dell does not recommend the use of file shares to host VM files due to the performance limitations. However, if you are using file shares, using a separate NIC is recommended.

Using Mount Points in Hyper-V

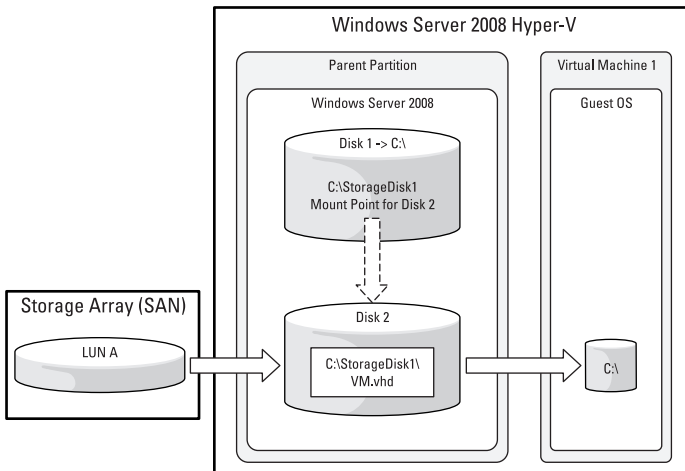
Failover Clustering in Windows Server 2008 allows a single cluster with up to 16 nodes. Therefore, you can have a single Hyper-V Host Cluster configuration across up to 16 servers. As shown in Figure B-1, one of the most common mechanisms to provision virtual disks to virtual machines is via VHD files created on a formatted volume on a storage disk mounted in the parent partition with a drive letter.

Figure B-1. Typical VHD Configuration on Storage Disk Assigned with a Drive Letter



In a Hyper-V Host Cluster configuration with many servers that host a large number of VHD-based virtual machines, there may not be sufficient number of drive letters (A-Z) in the parent partition to host the storage LUNs. To work around this drive letter limitation, Hyper-V allows the use of Mount Points to host storage LUNs.

Figure B-2. Typical VHD configuration on storage disk configured with a mount point



In Figure B-2, LUN A is not assigned a separate drive letter in the parent partition, but is instead provisioned in the parent partition using the mount point `C:\StorageDisk1`. For specific step-by-step information on configuring mount points, refer to Microsoft documentation at <http://support.microsoft.com/kb/947021>. Note that the use of mount points is an advanced technique to provision storage, and careful consideration should be provided before using it broadly.

In a Hyper-V Host Cluster environment, there are two additional requirements prior to successfully deploying HA VMs hosted on mount points:

- The Hyper-V and Failover Cluster Integration QFE (Microsoft KB951308 from <http://support.microsoft.com>) has to be applied on all Hyper-V nodes as well as systems running the **Failover Cluster Management** console. This is required for the **High Availability** wizard to detect that virtual machine files are on a volume that is using a mount point and not a drive letter.
- If any one of the Hyper-V nodes has a volume using a mount point, it is essential that all the Hyper-V nodes that are part of that Hyper-V Host Cluster have a folder with the exact name so that a placeholder exists when the storage disk fails over to any of the nodes. For example, in a configuration similar to that in the Figure B-2, all nodes should have a folder `C:\StorageDisk1`.

Hyper-V Host Clustering Data Template

The templates below provides guidance for data sheets that may be created to record information on the Hyper-V Host Cluster deployment. Note that the sample information within the tables are only for guidance.

Table C-1. Hyper-V Host Cluster Summary

Cluster Information	Cluster Solution
Hyper-V Host Cluster Name	
Hyper-V Host Cluster IP Address	
VM Workload Type	
Location	
Notes	

Table C-2. Hyper-V Host Cluster Nodes

Node Name	Service Tag	PowerEdge Server Model	Processor Info	Full or Core	Parent Partition / Public IP Address	Private IP Address

Table C-3. Hyper-V Host Cluster Network

Node Name	Physical NIC Info	Type	IP Address	Notes
	LOM #1	Host Management NIC & Cluster Public		
	LOM #2	Hyper-V Virtual Network		
	NIC #1	Hyper-V Virtual Network		
	NIC #2	iSCSI Storage NIC		
	NIC #3	Cluster Private		
	NIC #4	iSCSI Storage NIC		

Table C-4. Hyper-V Host Cluster HA VM Resource Group Information

VM Resource Group	VMs in Resource Group	Storage LUN Resources	Notes
Group-A-VMs	VM #1	Cluster Disk 1	
		Cluster Disk 2	
Group-B-VMs	VM #2	Cluster Disk 3	
		Cluster Disk 3	
		Cluster Disk 4	

References

- Dell Solutions for Windows Server 2008 Hyper-V at www.dell.com/hyper-v
 - *Dell Virtualization Solutions Advisor Tool*
 - *Dell Solutions Overview Guide for Microsoft Hyper-V*
 - *Dell Networking Solutions Guide for Microsoft Hyper-V*
 - *Dell Storage Solutions Guide for Microsoft Hyper-V*
 - *Dell Virtualization Reference Architecture for Microsoft Hyper-V*
- Dell High Availability Cluster Solutions documentation at <http://docs.us.dell.com/support/edocs/systems/clusters>
 - *HA Cluster Solutions featuring Dell | EMC Fibre Channel Storage*
 - *HA Cluster Solutions featuring Dell PowerVault SAS/SATA Storage*
 - *HA Cluster Solutions featuring Dell PowerVault and Dell | EMC iSCSI Storage*
 - *HA Cluster Solutions featuring Dell | EqualLogic PS Series iSCSI Storage*
- *Failover Clustering in Windows Server 2008* at www.microsoft.com/windowsserver2008/failover-clusters.msp
- *Microsoft Step-by-Step Guide to Configure a Two-Node File Server Failover Cluster* from www.microsoft.com
- *Microsoft Step-by-Step Guide for Testing Hyper-V and Failover Clustering* document at www.microsoft.com
- *Hyper-V RTM Update - Microsoft KB950050* from <http://support.microsoft.com>
- *Hyper-V and Failover Cluster Integration QFE - Microsoft KB951308* from <http://support.microsoft.com>

Glossary

FAILOVER CLUSTERING CONSOLE — The Microsoft Management Console (MMC) used to manage the Failover Clustering feature (cluadmin.msc) in Windows Server 2008.

HA VM — A VM that has been configured as a Windows Failover Cluster resource.

HYPER-V CLUSTER NODES — Windows Server 2008 servers that have the Hyper-V role configured and are part of a failover cluster.

HYPER-V MANAGER CONSOLE — The Microsoft Management Console (MMC) manages the Hyper-V Virtualization Role on Windows Server 2008.

LUN (LOGICAL UNIT NUMBER) — In this solutions guide, a LUN refers to a unique disk that is provisioned to a server from a storage enclosure or array.

MPIO — Multi-Path IO provides fault-tolerance and performance enhancements by implementing more than one physical connection path from servers to storage.

PASSTHROUGH DISKS — Physical disks (local disks or LUNs provisioned from SAN) that are configured directly to VMs.

VM — A Hyper-V Virtual Machine.

