

# Planning for Hyper-V

---

A Dell PowerVault™ MD3200 and MD3200i Series of  
Arrays Technical White Paper

Dell



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

© 2010 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

*Dell*, the *DELL* logo, and the *DELL* badge, *PowerConnect*, and *PowerVault* are trademarks of Dell Inc. *Symantec* and the *SYMANTEC* logo are trademarks or registered trademarks of Symantec Corporation or its affiliates in the US and other countries. *Microsoft*, *Windows*, *Windows Server*, and *Active Directory* 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.

May 2010

## Contents

Planning for Hyper-V on the Dell PowerVault MD3200 and MD3200i Series of Storage Arrays .....	2
Disclaimer.....	2
Supported Hardware and Software .....	2
Planning for Hyper-V Virtualized Configurations.....	3
Pre-installation Requirements.....	3
Planning the SAS Direct-Attach Installation with the Dell PowerVault MD3200 and MD3220 Storage Arrays .....	4
Direct-Attach Storage with SAS Considerations.....	4
Planning the IP SAN Installation with the Dell PowerVault MD3200i and MD3220i Storage Arrays .....	5
IP SAN Considerations.....	5
IP SAN Network Infrastructure for Performance.....	7
IP SAN Optimization for Throughput and I/Os .....	9
Considerations for Storage and Network Adapters (MD3200 and MD3200i Storage Arrays) .....	10
Connecting the Dell PowerVault MD3200 Storage Array.....	10
Connecting the Dell PowerVault MD3200i Storage Array .....	10
MD3200i Load-Balance Policies .....	11
Additional Considerations for Virtual Storage .....	12
Considerations When Using Pass-Through Physical Disks or VHDs .....	13
Directly-Mapped iSCSI Virtual Disks to the VM (Dell PowerVault MD3200i Storage Array) .....	14
Considerations for the VM Child Partition .....	14
Clustering with Hyper-V.....	14
Contact Information .....	17
Appendix.....	17

## Figures

Figure 1. Dell PowerVault MD3200 Storage Array Configuration.....	4
Figure 2. Dell PowerVault MD3200i Storage Array (IP SAN) .....	6
Figure 3. Dell PowerVault MD3200 Storage System (SAS) Cabling Configuration.....	10
Figure 4. Dell PowerVault MD3200i Storage Array Network Configuration.....	11
Figure 5. Virtual Storage Configuration.....	12
Figure 6. MD3200 SAS Storage Array Configuration in a Clustered Environment .....	16
Figure 7. MD3200i iSCSI Storage Array Configuration in a Clustered Environment .....	16

# Planning for Hyper-V on the Dell PowerVault™ MD3200 and MD3200i Series of Storage Arrays

This document helps you plan to implement Microsoft Hyper-V technology with the Dell PowerVault™ MD3200 and MD3200i series of storage arrays.

Hyper-V provides the software infrastructure and management tools in the Windows Server 2008 operating system to create and manage a virtualized server computing environment. The configuration in this document uses the Windows Server 2008 Release 2 operating system with the following enhancements to Hyper-V:

- Live Migration - A transparent move of virtual machines (VMs) from one cluster-node to another with no loss of VM connectivity between nodes. Live Migration uses clustered shared volumes (CSVs).
- Dynamic virtual machine storage - Support for hot plug-in and hot removal of storage.
- Enhanced processor support - Up to 32 physical cores.
- Enhanced networking support - Support for jumbo frames on VMs.

For information about how to deploy Hyper-V, refer to *Implementing Hyper-V on the Dell PowerVault MD3200 and MD3200i Series of Storage Arrays*. For information about backup and recovery options, refer to *Data Protection on the Dell PowerVault MD3200 and MD3200i Series of Storage Arrays in a Hyper-V Environment*.

This document provides useful information for system administrators and IT planners or designers who are in the process of considering or planning to use Hyper-V in their environment.

## *Disclaimer*

The steps in this document were performed on a test configuration. Live data was not used. This document implies no warranties. These instructions are not guaranteed to work in every situation. For the latest information about Dell IP SAN storage arrays, refer to the Dell storage array installation guides. Go to this website:

<http://www.dell.com>

## Supported Hardware and Software

Microsoft Hyper-V requires specific hardware. Refer to the Microsoft Windows Server catalog for Hyper-V to find out if your systems qualify. Go to this website:

<http://go.microsoft.com/fwlink/?LinkId=111228>

You also can check the Dell-supported hardware for Microsoft Hyper-V contained in the *Dell Solutions Overview Guide for Microsoft Hyper-V*. Go to this website:

<http://support.dell.com/support/edocs/software/HyperV/en>

If you need additional information, go to this website:

<http://www.dell.com/hyperv>

Both the Dell PowerVault MD3200 and MD3200i storage arrays are supported with the Microsoft Windows Server 2008 Release 2 operating system running the Hyper-V role.

## Planning for Hyper-V Virtualized Configurations

Planning for Hyper-V virtualized configurations includes these tasks:

- Making sure that pre-installation requirements are met
- Configuring the network adapter
- Deciding how to assign physical drives to VMs
- Deciding when to use pass-through physical disks or virtual hard drives (VHDs)
- Planning child partitions
- Deciding how to cluster

### Pre-installation Requirements

To install and use the Hyper-V role, make sure you have the following processing features on your system:

- An x64-based processor.

Hyper-V is not available for 32-bit systems or Itanium-based systems.

You can download the AMD utility for Hyper-V compatibility from this website:

- <http://support.amd.com/us/Pages/dynamicDetails.aspx?ListID=c5cd2c08-1432-4756-aafa-4d9dc646342f&ItemID=177>
- Hardware-assisted virtualization.

Processors with Intel Virtualization Technology (Intel VT)

AMD Virtualization Technology (AMD-V)

- Hardware-enforced data execution prevention (DEP) must be available and enabled.

The Intel XD bit must be enabled (execution disable bit).

The AMD NX bit must be enabled (no execute bit).

Install the updates that you need. To receive a current list of all of the updates, go to this website:

<http://technet.microsoft.com/en-us/library/dd430893.aspx>

## Planning the SAS Direct-Attach Installation with the Dell PowerVault™ MD3200 and MD3220 Storage Arrays

The Dell PowerVault™ MD3200 storage array, with SAS host interface connectivity, can connect with up to four servers for high availability. The configuration that you choose depends on the number of servers you have, whether you are connecting to servers or a cluster-node, and the level of data redundancy that your configuration requires.

### Direct-Attach Storage with SAS Considerations

#### Redundancy

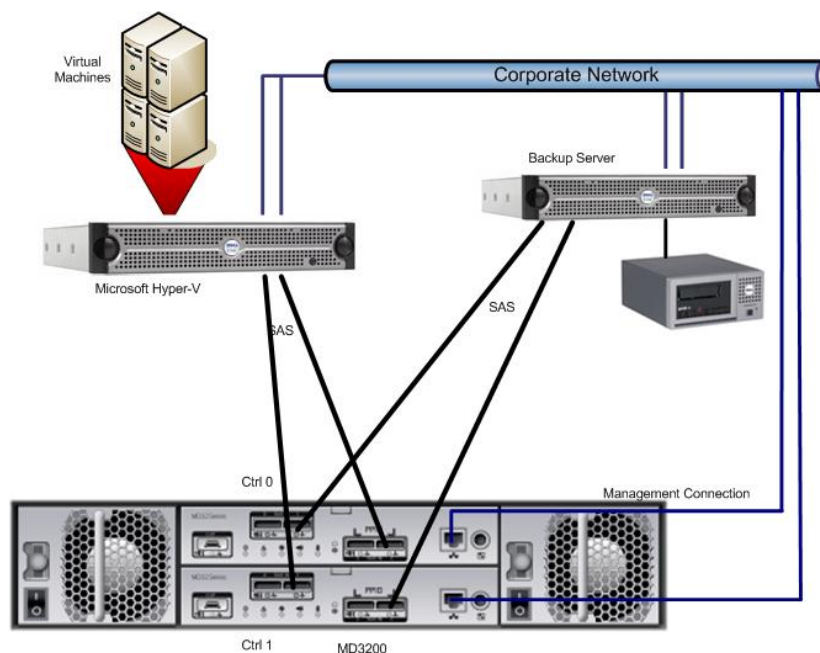
Redundancy is established by installing separate data paths between the server and the storage array, in which each path connects to the different RAID controller modules. Redundancy protects the server from losing access to data in the event of data path failure because both RAID controllers can access all of the disks in the storage array. Figure 1 shows redundant, high-availability cabling configuration for two hosts. This configuration is a best practice for Microsoft Hyper-V installations.

---

**Note** Use two dual-ported 6Gb host bus adapters (HBAs), and connect to one port from each of these HBAs to improve availability. With this configuration, a cable or a SAS HBA could fail, and the access to the data is still guaranteed.

---

**Figure 1.** Dell PowerVault MD3200 Storage Array Configuration



#### RAID Level

Choose an appropriate RAID level based on your applications. RAID 1 or higher provides some level of redundancy that is useful in the case of failed physical disks. Each RAID level works best with certain applications. Consider the RAID level when you configure the Dell PowerVault MD3200 storage array.

## Power

Connect the power supplies for the Dell PowerVault MD3200 storage array to separate power sources. Doing so ensures that even if one component fails due to a power issue, the alternative path continues to work.

# Planning the IP SAN Installation with the Dell PowerVault™ MD3200i and MD3220i Storage Arrays

You can implement an IP storage area network (SAN) in various ways based on business requirements, available resources, and intended application. For instance, one important but easily overlooked item that can improve the manageability of your IP SAN implementation is to assign a consistent and representative naming scheme to the storage arrays. This assignment is especially useful if the SAN has more than one storage array attached to it. You can use the “blink array” feature of the Modular Disk Storage Manager to correctly identify each storage array physically.

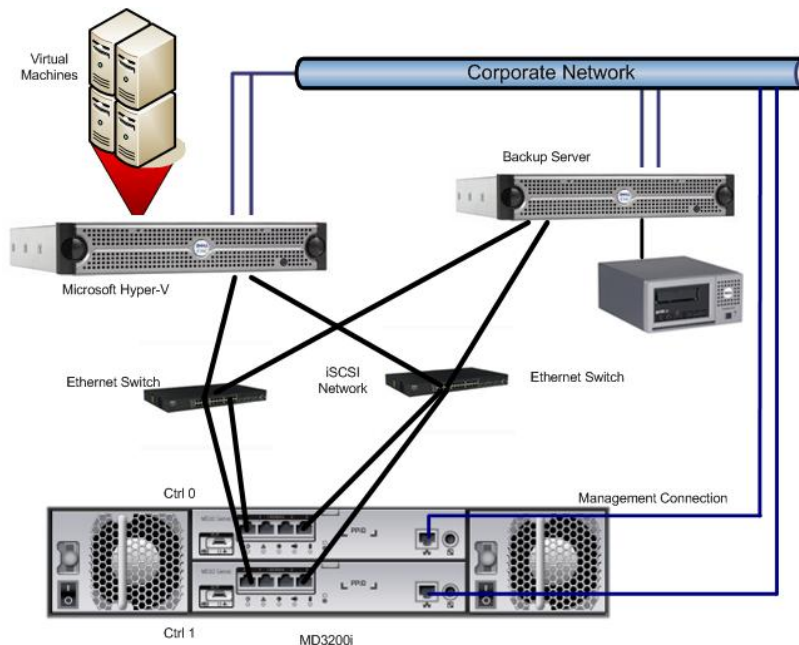
This section contains some general implementation guidelines; however, these guidelines might not benefit all applications.

## IP SAN Considerations

### Redundancy

Redundancy requires having a second set of hardware and communication paths so that if one piece of hardware on one path breaks down, a second path can be used. In an IP SAN, a second controller in the storage array can perform redundancy, and redundancy is available when you use two different switches for the iSCSI network. Figure 2 shows such a configuration.

**Figure 2.** Dell PowerVault MD3200i Storage Array (IP SAN)



### Physical Network Infrastructure

A fully redundant IP SAN is characterized by multiple physical independent iSCSI data paths between the servers and the storage array. Make sure that each data path is on a separate subnet.

### iSCSI Sessions

In the case of an iSCSI target, such as the Dell PowerVault MD3200i storage array, establish multiple sessions to the storage subsystem from each server. Set up one session per port from each of the network cards to each RAID controller module. This method allows one session to restart if a link goes down and does not affect any of the other links.

### RAID Level

Choose an appropriate RAID level based on your applications. RAID 1 or higher RAID levels provide some level of redundancy that is useful in the case of failed physical disks. Each RAID level works best with certain applications, so consider the RAID level when you configure the Dell PowerVault MD3200i storage array.

### Power

Connect the MD3200i power supplies to separate power sources. These connections ensure that if one component fails due to a power issue, the alternative path continues to work.

### Security

To secure data across an IP SAN, implement a physically independent network for the iSCSI data traffic, and isolate it from other Ethernet traffic. Having an isolated network dedicated to storage traffic



means that the storage array avoids network congestion with other non-storage traffic that is using the network.

- **CHAP**  
To have secure access between your host and storage array, enable target and mutual CHAP authentication on the hosts and storage arrays. Follow standard CHAP password guidelines for best security. Set a strong password that meets standard IT guidelines on all of the devices in your IP SAN.
- **Virtual LAN (VLAN)**  
If you want to physically isolate your IP SAN data traffic from general network traffic, use VLANs. Turn on VLAN tagging on the Dell PowerVault MD3200i storage array. A port can either transmit all tagged IP packets or all non-tagged IP packets.

---

**Note:** You must enable a VLAN throughout the entire iSCSI SAN from the network interface cards (NICs) and switches to the iSCSI ports; otherwise, behavior might be inconsistent. To simplify and avoid troubleshooting at deployment, make sure that the NICs, switches, and storage array are fully operational before you enable the VLAN feature for the entire solution.

---

## IP SAN Network Infrastructure for Performance

This section describes some general implementation guidelines for performance. These guidelines might not apply to all applications. Review your application requirements prior to implementing these guidelines.

### General Network Practices

Make sure the category rating for the cables used is gigabit-Ethernet-compliant (CAT5e or CAT6). Design your network to have the least amount of hops between the storage arrays and the servers. This design greatly reduces the failure points, it simplifies the manageability, and it reduces the latency and complexity of your network architecture (particularly in the area of redundancy).

Use managed switches because they provide advance features to help you optimize and maintain your network for your application. Only use auto-negotiation because gigabit Ethernet networks are designed to always have auto-negotiation enabled. If a particular application requires a specific speed or duplex mode, you can accomplish this by changing the advertisement options of the switch.

## Spanning Tree Protocol

Disable the spanning-tree protocol (STP) on the switch ports that connect end nodes (iSCSI initiators and SAN interfaces). If you still decide to enable STP on the switch ports, turn on the STP FastPort feature on the ports to allow immediate transition of the ports into a forwarding state.

---

**Note:** FastPort immediately transitions the port into the STP Forwarding mode upon linkup. Because the port still participates in STP and it is to be a part of the loop, it eventually transitions into the STP Blocking mode.

---

---

**Note:** PowerConnect switches default to Rapid Spanning Tree Protocol (RSTP), which is an evolution in STP that provides for faster spanning tree convergence and is preferable to STP when available.

---

---

**Note:** Use STP for a single cable connection between switches or for trunking for multiple cable connections between switches.

---

## TCP Congestion Avoidance

TCP congestion avoidance is an end-to-end flow control protocol that limits the amount of data sent between a TCP sender and a TCP transmitter. This protocol uses a sliding window to size the data being sent to the TCP receiver. This protocol starts with a small segment size and keeps increasing with each acknowledged segment sent until a segment is dropped. After a segment is dropped, TCP starts this avoidance over again.

## Ether Flow Control

It is best practice to enable flow control on the switch ports that handle iSCSI traffic. In addition, if a server is using a software iSCSI initiator-and-NIC combination to handle iSCSI traffic, enable flow control on the NICs to obtain a performance benefit. In SAN configurations in which many servers (initiators) are communicating with storage devices, if senders transmit data simultaneously, the data transmission might exceed the throughput capacity of the receiver. When this event occurs, the receiver might drop packets, which forces the senders to retransmit the data after a delay. Although this event does not result in any loss of data, latency increases because of the re-transmissions and I/O performance degrades. By enabling flow control, you can avoid the imbalance in the network traffic between the devices that send network traffic and the devices that receive the traffic.

---

**Note:** PowerConnect switches default to flow control being off. The Dell PowerVault MD3200i storage array auto-configures to the switch when the flow control is turned on.

---

## Unicast Storm Control

A traffic “storm” occurs when a large outpouring of IP packets creates excessive network traffic that degrades network performance. Many switches have traffic storm control features that prevent ports from being disrupted by broadcast, multicast, or unicast traffic storms on physical interfaces. These features typically work by discarding network packets when the traffic on an interface reaches a

percentage of the overall load (usually 80 percent, by default). Because iSCSI traffic is unicast traffic and can typically use the entire link, disable unicast storm control on switches that handle iSCSI traffic. However, the use of broadcast and multicast storm control is encouraged. Refer to your switch documentation for information on disabling unicast storm control.

### Jumbo Frames

Most enterprise gigabit-Ethernet equipment provides some support for jumbo frames. Enabling jumbo frames accelerates iSCSI performance by about 5 percent and reduces server CPU utilization by 2 percent to 3 percent with standard or smarter NICs. Enable jumbo frames on the switch ports that handle iSCSI traffic. In addition, if a server is using a software iSCSI initiator-and-NIC combination to handle iSCSI traffic, you must also enable jumbo frames on the NICs to obtain the performance benefit and ensure consistent behavior.

---

**Note:** You must enable jumbo frames throughout the entire IP SAN from the NICs, switches, and array ports; otherwise, behavior might be inconsistent. To simplify troubleshooting initial deployments, make sure that the NICs, switches, and storage arrays are fully operational before you enable jumbo frames.

---

## IP SAN Optimization for Throughput and I/Os

When designing your IP SAN, review the various factors in your network and the applications that you are using. Consider these general rules when you design your IP SAN.

To maximize the data throughput of your storage array, all data ports need to be used. If your application is I/O intensive, use an iSCSI offload NIC. Consider manually balancing your virtual disk ownership so that no single controller is processing an excessive amount of I/O relative to the other controller. The Dell PowerVault MD3200i storage system supports active/active controllers, with each controller being able to simultaneously process I/O. The asymmetric design of the controllers means that a virtual disk (logical unit number or LUN) is owned by a controller, and all I/O access to the virtual disk is only possible through the owning controller. To take advantage of both the controllers for I/O access, distribute virtual disks among the controllers and modify them to balance I/O access so as to balance controller utilization.

### Layer 3 Optimization

Differentiated Services (DiffServ) provide a good method for managing your network traffic. Some switches have a proprietary implementation of DiffServ called Quality of Service (QoS). DiffServ uses the Differentiated Services Code Point (DSCP) to distinguish between service levels of each IP connection. These service-level agreements are on a per-hop basis (PHB). As such, within the internal corporate network, traffic flows can be predictable, but after a wide area network (WAN) link leaves the company, the service-level agreements are no longer valid. Four levels are typically used with DiffServ:

- Default PHB - Typically this level is best-effort traffic.
- Expedited forwarding (EF) PHB - Good for low-loss, low-latency traffic.
- Assured forwarding (AF) - Behavior group.
- Class selector PHBs - Maintain backward compatibility with the IP precedence field.

To choose the service level to use, examine the needs of the applications connected to the storage array. For instance, if you have your servers set up to iSCSI boot, or are using virtualization to “hide” the storage array (and the guest OS is booting off a C: drive that is actually on the storage array), select EF because the data must get to the server and, if there is latency, the server locks up. On the other hand, if you might want all of your traffic coming in from the worldwide web, set to the lowest possible class of AF so that the service level does not affect your critical data.


## Considerations for Storage and Network Adapters (PowerVault™ MD3200 and MD3200i Storage Arrays)

When planning for the network configuration, make sure that you have at least two network adapters assigned for your virtual machine (VM) traffic. If you are planning to add MS-Exchange, MS-SQL, MS-SharePoint, or a similar configuration to your VMs, use two dedicated network connections for each of those applications. For example, if you have a file server, MS-Exchange, and a SQL server running on VMs, you should have six network connections (two for MS-Exchange, two for the SQL server, one for the file server, and one for managing the Hyper-V server). For clustering, have two additional network connections for the cluster heartbeat.

### Connecting the Dell PowerVault MD3200 Storage Array

The connection to the Dell PowerVault MD3200 storage array is SAS based. **Error! Reference source not found.** shows how the SAS storage cabling is configured.

**Figure 3.** Dell PowerVault MD3200 Storage System (SAS) Cabling Configuration

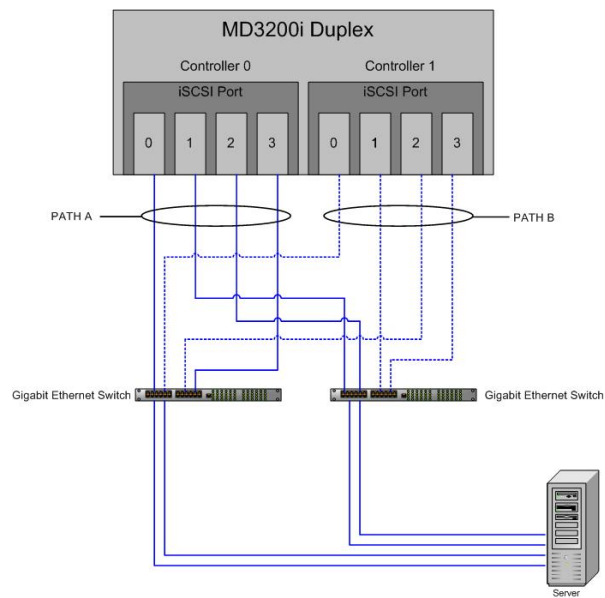


### Connecting the Dell PowerVault MD3200i Storage Array

For the iSCSI-based Dell PowerVault MD3200i storage array, you must have two NICs for I/O to the storage array for each server. Set up the management ports (out-of-band) for managing the MD3200i storage array because they are required for the initial configuration. For best performance when

planning the network connection to the Dell PowerVault MD3200i storage array, make sure that the cabling looks similar to **Error! Reference source not found.** (this configuration uses maximum performance by using all four ports per controller; however, a minimum of two connections is recommended).

**Figure 4.** Dell PowerVault MD3200i Storage Array Network Configuration



The Module Disk Storage Manager comes with the device-specific module (DSM) for Microsoft multipath I/O (MPIO). MPIO provides load-balancing between the two ports of one RAID controller module. For best performance, make sure that load-balancing is enabled and active. Load-balancing lets you use up to 4 Gb/s for each virtual disk on the Dell PowerVault MD3200i storage array.

#### *MD3200i Load-Balance Policies*

You can change the load-balance policy by using the Windows 2008 Disk Management option. From the server desktop, right-click the **My Computer** icon, and select **Manage** to open the Computer Management dialog. Click **Disk Management** to show the list of virtual disks that are attached to the server. Right-click the virtual disk on which you would like to set the load-balance policy, and select **Properties**. From the **MPIO** tab, select the load-balance policy that you want to set for this virtual disk

You can choose one of these load-balance policies to optimize I/O performance.

#### **Round Robin with Subset**

The round-robin with subset I/O load-balance policy routes I/O requests, in rotation, to each available data path to the RAID controller module that owns the virtual disks. This policy treats all data paths to the RAID controller module that owns the virtual disk equally for I/O activity. Data paths to the

secondary RAID controller module are ignored until the ownership changes. The basic assumption for the round-robin policy is that the data paths are equal.

#### Least Queue Depth with Subset

The least queue depth with subset policy is also known as the least I/Os or least requests, policy. This policy routes the next I/O request to a data path that has the least outstanding I/O requests queued. For this policy, an I/O request is simply a command in the queue. The type of command or the number of blocks that are associated with the command are not considered. The least queue depth with subset policy treats large block requests and small block requests equally. The data path selected is the path group of the RAID controller module that owns the virtual disk.

#### Least Path Weight with Subset

The least path weight with subset policy assigns a weight factor to each data path to a virtual disk. An I/O request is routed to the data path with the lowest weight value to the RAID controller module that owns the virtual disk. If more than one data path to the virtual disk has the same weight value, the round-robin with subset load-balance policy is used to route I/O requests between data paths with the same weight value.

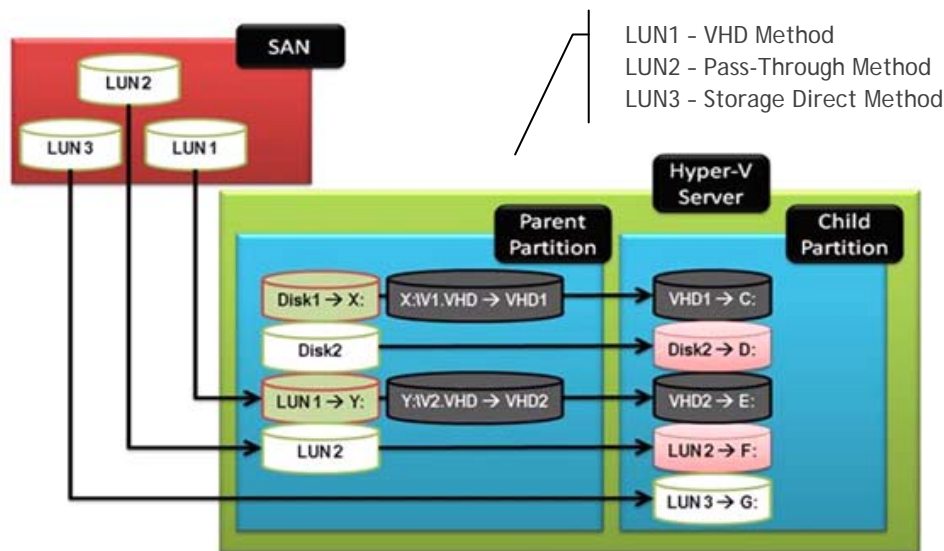
## Additional Considerations for Virtual Storage

You can assign physical storage to a virtual machine in various ways. For the Dell PowerVault™ MD3200 or MD3200i storage array, you have two options:

- Assign the virtual disks to the parent server, and create a file system and a virtual hard drive (VHD) for the virtual machine.
- Assign an unformatted physical disk directly to the VM (pass-through physical disk).

With the Dell PowerVault MD3200i storage array, in addition to the previously mentioned two options, you have a third option: You can assign a physical disk directly to the VM by installing the iSCSI software initiator on the VM as you would on a physical server.

### **Figure 5.** Virtual Storage Configuration



VHDs have several different options:

- Dynamic expansion - The VHD file grows as data is stored to the physical disk, with a maximum size not exceeding the size specified during creation. When using the New Virtual Machine Wizard in the Hyper-V Manager, select **Dynamically Expanding** when you create a new VM.
- Fixed size - The VHD file uses the full amount of space specified during creation.
- Differencing - The VHD file exists as a child disk in a parent/child relationship with an existing VHD. The parent VHD is left intact, and changes that occur are stored on the child (differencing) VHD.

---

**Note:** For the best performance in production environments, use fixed-size VHDs because they eliminate the overhead created by expanding the VHD in a normal production environment.

---

Size each LUN to support the target VHD. To avoid warning messages generated by the parent partition about the physical disk running out of capacity, allocate an additional 200 MB to the LUN, and leave this additional 200 MB as empty space on the virtual disk.

## Considerations When Using Pass-Through Physical Disks or VHDs

Keep the following considerations in mind when you use pass-through physical disks or VHDs:

- VHDs support the ability to export the drive to another parent server.
- VHDs support VM checkpoints.
- Pass-through physical disks offer slightly higher performance than VHDs and support virtual disk sizes larger than 2 TB. However, pass-through physical disks do not support VSS-based backup or checkpoints. VSS-based backups with hardware providers work as long as the backups are initiated from the parent. Inside the VM, only software-based VSS backup is supported.
- You can use both VHDs and pass-through physical disks in clustered environments.

- If you are booting from a pass-through physical disk, you must store the VM files in a different location.
- For every VM, you must create a separate pass-through physical disk. You cannot share a pass-through physical disk for two VMs.
- The physical layout of the pass-through physical disk is the same as for any physical server.

---

**Note:** For the highest performance, use pass-through physical disks but only as data devices. If you need the best flexibility, assign fixed-size VHDs.

---

For more information about how to maximize performance for different applications on the Dell PowerVault MD3200 and MD3200i storage arrays, go to this website:

[http://www.dell.com/downloads/global/products/pvaul/en/pvault\\_MD3200\\_whitepaper.pdf](http://www.dell.com/downloads/global/products/pvaul/en/pvault_MD3200_whitepaper.pdf)

---

**Note:** The configuration of your physical disk layout for performance on the Dell PowerVault MD3200 or MD3200i storage array depends on the application requirements of your VM (child partition), not on the Hyper-V parent server.

---

## Directly-Mapped iSCSI Virtual Disks to the VM (Dell PowerVault MD3200i Storage Array)

You can easily transfer the virtual disk from a physical server to the VM. However, transferring the virtual disk reduces the number of servers that are connected to the Dell PowerVault MD3200i storage array. Every VM uses exactly one server partition. VSS-based backup from the parent does not include the directly mapped iSCSI LUNs of the child partition. Booting also is not possible unless you use third-party software. The Dell PowerVault MD3200i storage array requires only a software initiator; the iSCSI HBA is not supported with the Dell PowerVault MD3200i storage array. Only software initiators are supported with the Dell PowerVault MD3200i storage array.

---

**Note:** You can use direct mapping for data devices that were originally used on a physical server or for data devices that are being moved back to a physical server.

---

## Considerations for the VM Child Partition

Usually, when you create VHDs, you can store your VM configuration files in the same location as the VHDs. If you plan to use multiple VMs, you can create one virtual disk on your Dell PowerVault MD3200 or MD3000i storage array for storing your configuration files and assign the data physical disks separately. Depending on the performance required for your different VMs, you might use one MD3200 or MD3000i virtual disk for multiple VM data disks. When performing capacity calculations, add approximately 200 MB to the VM configuration data for each VM. When configuring the VM, booting is supported only on virtual IDE adapters, not on the virtual SCSI adapter. However, you can assign an MD3200 or an MD3000i virtual disk as an IDE device to the VM, and performance is not affected.

If you plan to use pass-through physical disks as boot disks and data disks, you must store your VM configuration files in a different location. This location can be the local physical disk of your server, or an MD3200 or MD3000i virtual disk, which is available to the parent and formatted with NTFS. You can assign a pass-through device only after creating the VM. To assign a pass-through device, select **Assign Hard Drive Later** during the VM creation process. After the VM is created, change the VM options to **Assign the Pass-Through Device**.



---

**Note:** The pass-through physical disk must have a signature and be in an Offline state on the parent so that you can assign the physical disk to a VM.

---

## Clustering with Hyper-V

If you want to use clustering for Hyper-V, carefully observe the following requirements.

- Use identical server hardware to eliminate issues regarding failover.
- Avoid having a single point-of-failure in your communications network. Connect cluster-nodes through a network that uses redundant switches or similar hardware.

---

**Note:** You can connect cluster-nodes with a single network. However, it is best to use a redundant communications network.

---

- All cluster disks must be basic disks. Dynamic disks are not supported.

### PowerVault™ MD3200 storage array-specific cluster requirements:

- Only use two dual-ported 6Gb SAS adapters for each server, and connect one port of each adapter to the Dell PowerVault MD3200 storage array in clustered environments to achieve the highest level of redundancy.

### PowerVault™ MD3200i storage array-specific cluster requirements:

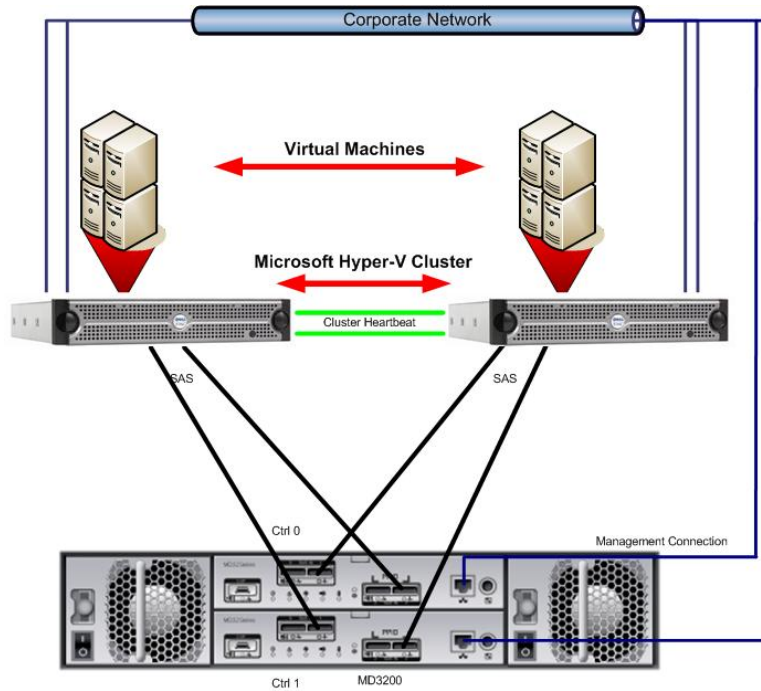
- The storage configuration must be redundant. Use two or more network connections for iSCSI traffic to the Dell PowerVault MD3200i storage array.
- You must have dedicated iSCSI network adapters for iSCSI environments. You cannot share communication and iSCSI traffic on the same adapter in clustered configurations.

All cluster-nodes must be able to communicate with the same virtual disks on the Dell PowerVault MD3200 or MD3200i storage array. You can use storage partitioning to create a server group with the cluster-nodes and assign the necessary virtual disks to them. Make sure that only the cluster-nodes have access to those virtual disks and that other servers do not have access to those virtual disks.


When planning, create an additional virtual disk for the quorum disk. Make this virtual disk at least 1 GB in size. This disk contains the quorum information and is used as the witness disk. Format the quorum virtual disk with NTFS.

If you plan to use pass-through devices on your VMs, you must make sure that the VM configuration files are also stored on the shared storage and that the shared storage is available to all cluster nodes. You can create an additional virtual disk on the Dell PowerVault MD3200 or MD3200i storage array or increase the size of the quorum virtual disk to store the configuration files there. You also might create an additional virtual disk with approximately 300 MB to 500 MB allocated for each VM. For example, 10 VMs equal a 5GB disk size.

**Figure 6.** MD3200 SAS Storage Array Configuration in a Clustered Environment



**Figure 7.** MD3200i iSCSI Storage Array Configuration in a Clustered Environment



## Contact Information

Go to this website:

<http://www.dell.com/>

Or call 1-800-WWW-DELL.

## Appendix

### *Appendix A: System Troubleshooting*

For information about troubleshooting MD3200 storage arrays, refer to the “Troubleshooting Problems” chapter of the *Dell™ PowerVault™ Modular Disk Storage Manager User’s Guide*. Go to this website:

<http://support.dell.com/support/edocs/systems/MD3200/en/index.htm>

### *Appendix B: References*

Cisco, undated. *Flooding Control*

[http://www.cisco.com/univercd/cc/td/doc/product/lan/c2900xl/29\\_35xu/olhelp/floodctl.htm](http://www.cisco.com/univercd/cc/td/doc/product/lan/c2900xl/29_35xu/olhelp/floodctl.htm)

Dell, undated. *Dell™ PowerVault™ Modular Disk Storage Manager CLI Guide*

<http://support.dell.com/support/edocs/systems/MD3200/en/index.htm>.

Dell, 2007. *Manuals: Dell™ PowerVault™ MD3200*

<http://support.dell.com/support/edocs/systems/MD3200/en/index.htm>

Dell, 2007. *PowerVault MD3200i SAN Array for Storage Consolidation*

[http://www.dell.com/content/products/productdetails.aspx/pvaul\\_MD3200i?c=us&1=en&s=bsd&cs=04](http://www.dell.com/content/products/productdetails.aspx/pvaul_MD3200i?c=us&1=en&s=bsd&cs=04)

Dell, 2007. *Using iSCSI: Dell™ PowerVault™ Modular Disk Storage Manager User’s Guide*

<http://support.dell.com/support/edocs/systems/MD3200/en/UG/HTML/iscsi.htm>

Microsoft, 2007. *How to Align Exchange I/O with Storage Track Boundaries*

<http://technet.microsoft.com/en-us/library/aa998219.aspx>

Microsoft, 2007. *Planning Storage Configurations for Exchange 2007*

<http://technet.microsoft.com/en-us/library/bb124518.aspx>

Microsoft, 2008. *Microsoft Storage Technologies - iSCSI*

<http://www.microsoft.com/windowsserver2003/technologies/storage/iscsi/default.mspx> (FAQs and more information on Microsoft iSCSI)

Planning for Hyper-V: PowerVault™ MD3200 and MD3200i Series of Arrays

Microsoft Download Center, 2007. *Microsoft iSNS Server 3.0*

<http://www.microsoft.com/downloads/details.aspx?familyid=0dbc4af5-9410-4080-a545-f90b45650e20&displaylang=en>

Microsoft Exchange Team Blog, 2007. *Exchange 2007 Mailbox Server Role Storage Requirements Calculator*

<http://msexchangeteam.com/archive/2007/01/15/432207.aspx>

### *Appendix C: Terminology*

**Bandwidth** - The amount of data that the storage array can process over time. Bandwidth is measured in megabytes per second (MB/s).

**Child partition** - A partition that is associated with the parent partition. The child partition is subordinate to the parent partition.

**Initiator** - The system component that originates an input/output (I/O) command over an I/O bus or network. I/O adapters, network interface cards, and intelligent RAID controller module device I/O bus control application specific integrated circuits (ASICs) are typical initiators. A peripheral device is a target. (*The Dictionary of Storage Networking Terminology*)

**Latency** - The interval of time between submitting a request and receiving a response.

**MPIO (Multipath I/O)** - In computer storage, an arrangement whereby more than one logical path exists between the central processing unit (CPU) in a computer system and its storage devices. This path is routed through the buses and bridge devices that connect the CPU and its storage devices. If one RAID controller module, port, or switch fails, the operating system can route I/O through the remaining RAID controller module so that work can continue.

**NIC (network interface card)** - An adapter that connects an intelligent device to a network. The common terminology is a network interface card or an Ethernet NIC network.

**Node** - An addressable entity connected to an input/output (I/O) bus or network and used primarily to refer to computers, storage devices, and storage arrays. The component of a node that connects to the bus or network is a port. (*The Dictionary of Storage Networking Terminology*).

**Parent Partition** - The primary partition. The parent partition creates links to child partitions, and has access to those child partitions.

**Pass-through Physical Disk** - An unformatted physical disk directly assigned to the virtual machine.

**Portal** - A service that links initiators and targets when their IP addresses are not known.

**Quorum Virtual Disk** - A quorum virtual disk is a virtual disk on the shared storage array that stores the quorum information.

**Response time** - The interval of time between submitting a request and receiving a response.

**Session** - A group of iSCSI connections. As each connection is established, the session can negotiate its own unique parameters.

**Signature** - An embedded code that identifies the device. The signature is used when passing information between devices.

**Target** - The storage destination. In this document, the target is the MD3200i storage array.

**TCP/IP (Transmission Control Protocol/ Internet Protocol)** - A suite of protocols that includes Transmission Control Protocol (TCP), Internet Protocol (IP), User Datagram Protocol (UDP), and Internet Control Message Protocol (ICMP). These protocols form the basic set of communication protocols used on the Internet. (*The Dictionary of Storage Networking Terminology*)

**Throughput** – The number of individual I/Os the storage array can process over time. Throughput is measured in I/Os per second (IOPS).