

Realtime  
publishers

*The Shortcut Guide<sup>™</sup> To*



**Architecting  
iSCSI Storage for  
Microsoft Hyper-V**

*sponsored by*



*Greg Shields*

---

# Introduction to Realtime Publishers

---

by Don Jones, Series Editor

For several years now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We’ve made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book’s production expenses for the benefit of our readers.

Although we’ve always offered our publications to you for free, don’t think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you \$40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the “realtime” aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We’re an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I’m proud that we’ve produced so many quality books over the past years.

I want to extend an invitation to visit us at <http://nexus.realtimepublishers.com>, especially if you’ve received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you’re sure to find something that’s of interest to you—and it won’t cost you a thing. We hope you’ll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Introduction to Realtime Publishers..... i

Chapter 1: The Power of iSCSI in Microsoft Virtualization ..... 1

    The Goal for SAN Availability Is “No Nines” ..... 2

    Hyper-V Is Exceptionally Dependant on Storage..... 3

        VHD Attachment to VM..... 4

        Pass-Through Disks ..... 6

        iSCSI Direct Attachment..... 7

        VM-to-VM Clustering..... 9

        Host Boot from SAN..... 9

        Guest Boot from SAN..... 9

    VM Performance Depends on Storage Performance..... 10

        Network Contention ..... 11

        Connection Redundancy & Aggregation ..... 12

        Type and Rotation Speed of Drives..... 12

        Spindle Contention..... 13

        Connection Medium & Administrative Complexity ..... 13

iSCSI Makes Sense for Hyper-V Environments..... 15

## Copyright Statement

© 2009 Realtime Publishers. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (the "Materials") and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers its web site sponsors. In no event shall Realtime Publishers or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at [info@realtimepublishers.com](mailto:info@realtimepublishers.com).

[**Editor's Note:** This eBook was downloaded from Realtime Nexus—The Digital Library for IT Professionals. All leading technology eBooks and guides from Realtime Publishers can be found at <http://nexus.realtimepublishers.com>.]

# Chapter 1: The Power of iSCSI in Microsoft Virtualization

---

Virtualization is one of the hottest technologies to hit IT in years, with Microsoft's Hyper-V R2 release igniting those flames even further. Hyper-V arrives as a cost-effective virtualization solution that can be easily implemented by even the newest of technology generalists.

But while Hyper-V itself is a trivial implementation, ensuring its highest levels of redundancy, availability, and most importantly performance are not. Due to virtualization's heavy reliance on storage, two of the most critical decisions you will make in implementing Hyper-V are where and how you'll store your virtual machines (VMs).

Virtualization solutions such as Hyper-V enable many fantastic optimizations for the IT environment: VMs can be easily backed up and restored in whole, making affordable server restoration and disaster recovery possible. VM processing can be load balanced across any number of hosts, ensuring that you're getting the most value out of your server hardware dollars. VMs themselves can be rapidly deployed, snapshotted, and reconfigured as needed, to gain levels of operational agility never before seen in IT.

Yet at the same time virtualization also adds levels of complexity to the IT environment. Gone are the traditional notions of the physical server "chassis" and its independent connections to networks and storage. Replacing this old mindset are new approaches that leverage the network itself as the transmission medium for storage. With the entry of enterprise-worthy iSCSI solutions into the market, IT environments of all sizes can leverage the *very same network infrastructure they've built over time* to host their storage as well. This already-present network pervasiveness combined with the dynamic nature of virtualization makes iSCSI a perfect fit for your storage needs.

Correctly connecting all the pieces, however, is the challenge. To help, this guide digs deep into the decisions that environments large and small must consider. It looks at best practices for Hyper-V storage topologies and technologies, as well as cost and manageability implications for the solutions available on the market today. Both this and the following chapter will start by discussing the technical architectures required to create a highly-available Hyper-V infrastructure. In Chapter 2, you'll be impressed to discover just how many ways that redundancy can be inexpensively added to a Hyper-V environment using native tools alone.

If, like many, your storage experience is thus far limited to the disks you plug directly into your servers, you'll be surprised at the capabilities today's iSCSI solutions offer. Whereas Chapters 1 and 2 deal with the interconnections between server and storage, Chapter 3 focuses exclusively on capabilities within the storage itself. Supporting features such as automatic restriping, thin provisioning, and built-in replication, today's iSCSI storage provides enterprise features in a low-cost form factor.

Finally, no storage discussion is fully complete without a look at the affordable disaster recovery options made available by virtualizing. Chapter 4 discusses how iSCSI's backup, replication, and restore capabilities make disaster recovery solutions (and not just plans) a real possibility for everyone.

But before we delve into those topics, we first need to start with your SAN architecture itself. That architecture can arguably be the center of your entire IT infrastructure.

## The Goal for SAN Availability Is “No Nines”

It has been stated in the industry that “The goal for SAN availability is ‘no nines’ or 100% availability.” This is absolutely true in environments where data loss or non-availability have a recognizable impact on the bottom line. If your business loses thousands of dollars for every second its data is not available, you'd better have a storage system that never, ever goes down.

While such a goal could be laughable if it were applied to general-purpose operating systems (OSs) such as Microsoft Windows, 100% availability is not unheard of in the specialized hardware solutions that comprise today's SANs. No matter which company builds your SAN, nor which medium it uses to transfer data, its single-purpose mission means that multiple layers of redundancy can be built-in to its hardware:

- **Multiple power supplies** means that no single cable or power input loss can cause a failure.
- **Multiple and redundant connections** between servers and storage ensure that a connection loss can be survived.
- **Redundant pathing** through completely-isolated equipment further protects connection loss by providing an entirely separate path in the case of a downstream failure.
- **RAID configurations** ensure that the loss of a single disk drive will not cause the loss of an entire volume of data.
- **Advanced RAID configurations** further protect against drive loss by ensuring that even multiple, simultaneous drive failures will not impact availability.
- **Data striping across storage nodes** creates the ultimate protection by preserving availability even after the complete loss of SAN hardware.

All these redundant technologies are laid in place because a business' data is its most critical asset. Whether that data is contained within Microsoft Office documents or high-performance databases, any loss of data is fundamentally critical to a business' operations.

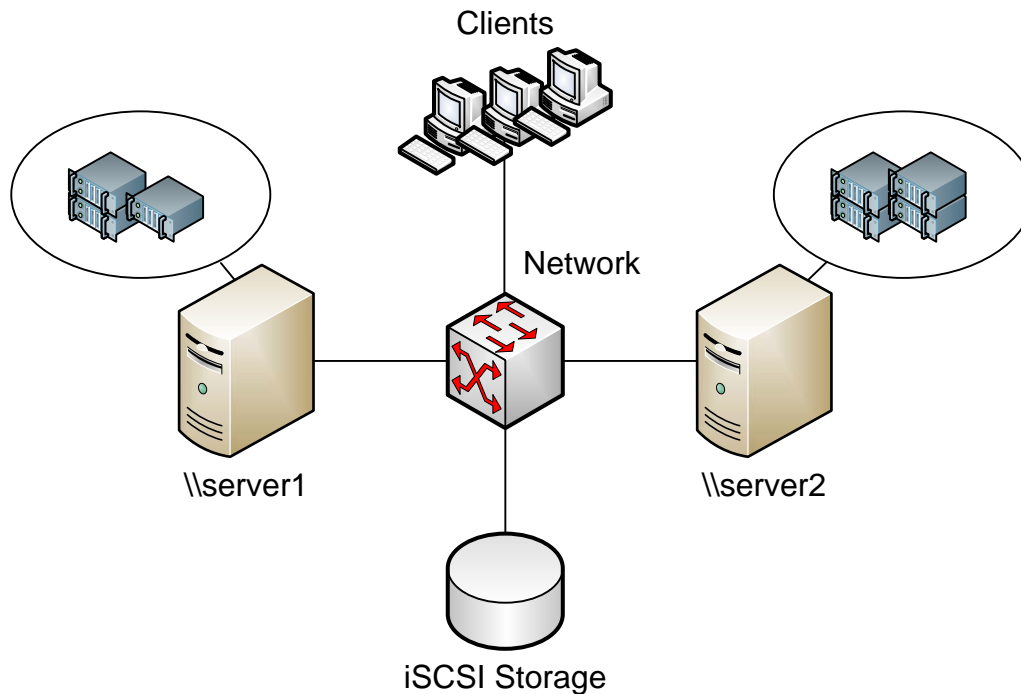
Yet a business' data is only one facet of the IT environment. That data is useless without the applications that work with it and create meaning out of its bits and bytes. In a traditional IT environment, those applications run atop individual physical servers, with OSs and applications often installed to their own local direct-attached storage. While your applications' data might sit within a highly-available SAN, the thousands of files that comprise each OS and its applications usually remain local.

With virtualization, everything changes. Moving that same environment to virtualization encapsulates each server's OS and its applications into a virtual disk. That virtual disk is then stored within the very same SAN infrastructure as your business data. As a result, making the move to virtualization effectively elevates your run-of-the-mill OS and application files *to the same criticality as your business data*.

## Hyper-V Is Exceptionally Dependant on Storage

Let's take a look at the multiple ways in which this new criticality occurs. Figure 1.1 depicts an extremely simplistic representation of a two-node Hyper-V cluster. In this cluster, each server connects via one or more interfaces to the environment's network infrastructure. Through that network, the VMs atop each server communicate with clients to provide their assigned services.

Important to recognize here is that high-availability in Hyper-V—like most virtualization platforms—requires some form of shared storage to exist across every host in the cluster.



**Figure 1.1: Highly-available Hyper-V at its simplest.**

That shared storage is the location where Hyper-V's VMs reside. This is the case because today's VM high-availability technologies never actually move the VM's disk file. Whether the transfer of ownership between two hosts occurs as a live migration with a running VM or a re-hosting after a physical host failure, the high-availability relocation of a VM *only moves the processing and not the storage of that VM.*

It is for this reason the storage component of any Hyper-V cluster is its most critical element. Every VM sits within that storage, every Hyper-V host connects to it, and all the processing of your data center's applications and data are now centralized onto that single device.

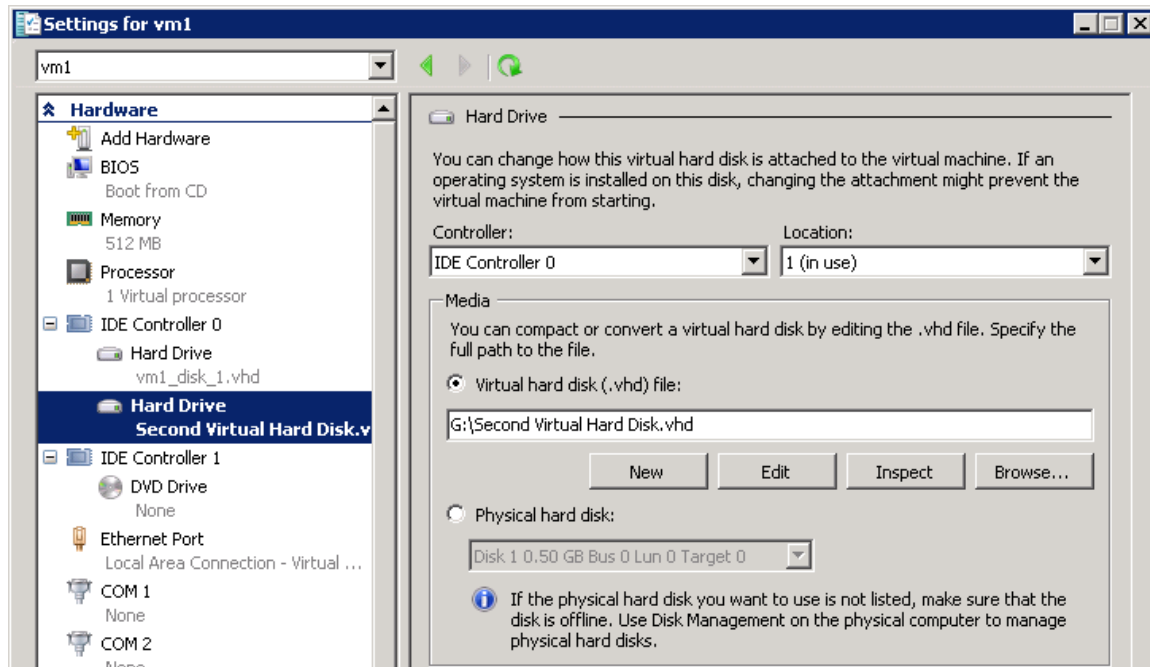
Yet this is only the simplest of ways in which a Hyper-V cluster interacts with its storage. Remember that iSCSI is in effect an encapsulation of traditional SCSI commands into network-routable packets. This encapsulation means that wherever your network exists, *so can your storage.* As a result, there are a number of additional ways in which virtual hosts and machines can connect to their needed storage. Let's take a look through a few that relate specifically to Hyper-V's VMs. You'll find that not all options for connecting VMs to storage are created alike.

### VHD Attachment to VM

Creating a new VM requires assigning its needed resources. Those resources include one or more virtual processors, a quantity of RAM, any peripheral connections, and the disk files that contain its data. Any created VM requires at a minimum a single virtual hard disk (VHD) to become its storage location.



Although a single VHD is the minimum, it is possible to attach additional VHDs to a VM either during its creation or at any point thereafter. Each newly-attached VHD becomes yet another drive on the VM. Figure 1.2 shows how a second VHD, stored at G:\Second Virtual Hard Disk.vhd, has been connected to the VM named `vm1`.



**Figure 1.2: Attaching a second VHD to an existing VM.**

Attached VHDs are useful because they retain the encapsulation of system files into their single .VHD file. This encapsulation makes them portable, enabling them to be disconnected from one VM and attached to another at any point. As VHDs, they can also be backed up as a single file using backup software that is installed to the Hyper-V host, making their single-file restore possible.

However, VHDs can be problematic when backup software requires direct access to disks for proper backups or individual file and folder restores. Also, some applications require an in-band and unfiltered SCSI connection to connected disks. These applications, while rare, will not work with attached VHD files. Lastly, VHDs can only be connected or disconnected when VMs are powered down, forcing any change to involve a period of downtime to the server.

VHDs can be created with a pre-allocated fixed size or can be configured to dynamically expand as data is added to the VM. All VHDs are limited to 2040GB (or just shy of 2TB) in size. Dynamically expanding VMs obviously reduces the initial amount of disk space consumed by the freshly-created VM. However, care must be taken when collocating multiple dynamically-expanding VHD files on a single volume, as the combination of each VHD's configured maximum size will often be greater than the maximum size of the volume itself. Proactive monitoring must be laid into place to watch for and alert on growth in the size of storage when dynamically-expanding VHDs are used.

The level of expected performance between fixed and dynamic VHD files is only slightly different when using Hyper-V R2, with fixed disks seeing a slightly increased level of performance over those created as dynamic. Dynamic VHD files incur an overhead during write operations that expand the VHD's size, causing a slight reduction in performance over fixed disks. Microsoft testing suggests that fixed VHDs see performance that is equal to native disk performance when run atop Hyper-V R2. Dynamic disks experience between 85% and 94% of native performance, depending on the type of write operations being done within the VM.

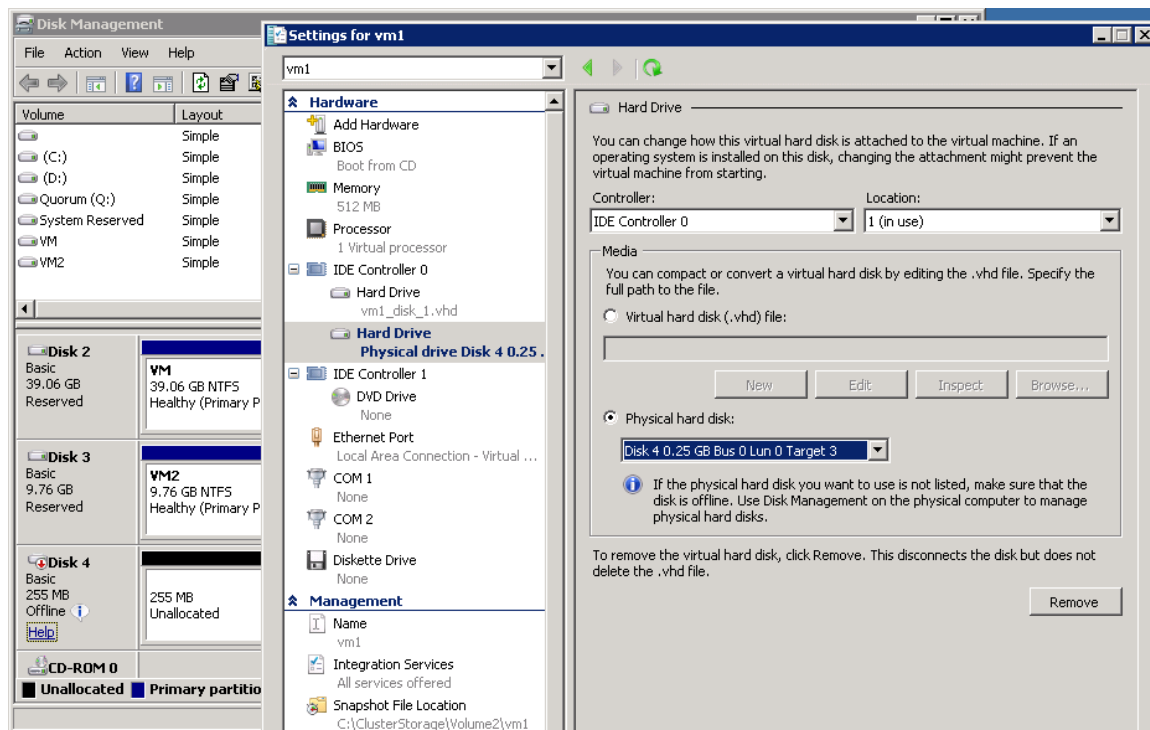
Your decision about whether to use fixed versus dynamic VHDs will depend on your need for slightly better performance versus your available quantity of storage. Consumed storage, however, does represent a cost. As you'll discover in Chapter 3, the capability for thin-provisioning VM storage often outweighs any slight improvements in performance.

### Pass-Through Disks

An alternative approach to pulling extra disks into a VM is through the creation of a pass-through disk. With this approach, an iSCSI disk is exposed to the Hyper-V host and then *passed through* from the host to a residing VM. By passing through the disk rather than encapsulating it into a VHD, its contents remain in their native format. This allows certain types of backup and other software to maintain direct access to the disk using native SCSI commands. As essentially raw mappings, pass-through disks also eliminate the 2040GB size limitation of VHDs, which can be a problem for very large file stores or databases.

Microsoft suggests that pass-through disks achieve levels of performance that are equivalent to connected VHD files. Pass-through disks can also be leveraged in clustered Hyper-V scenarios by creating the disk as a clustered resource after assigning it to a VM.

Figure 1.3 shows how a pass-through disk is created between a host and its residing VM. Here, as in the previous example, pass-through disks can only be attached to VMs that have been powered off. In this image, the host's Disk Management wizard has been displayed on the left with a 256MB offline disk attached via iSCSI and initialized by the host. Once initialized, the disk is taken offline and made available to the VM through its settings wizard on the right. There, the VM's second disk drive is attached to the passed-through hard drive, which is labeled *Disk 4 0.25 GB Bus 0 Lun 0 Target 3*.



**Figure 1.3: Creating a pass-through disk.**

Because they are not encapsulated into VHDs, pass-through disks cannot be snapshotted by Hyper-V. However, because the files reside on-disk in a native format, your storage solution may be able to complete the snapshot from its own perspective. This storage-level snapshot can enable advanced storage-level management functions such as replication, backup and restore, and volume-level cloning.

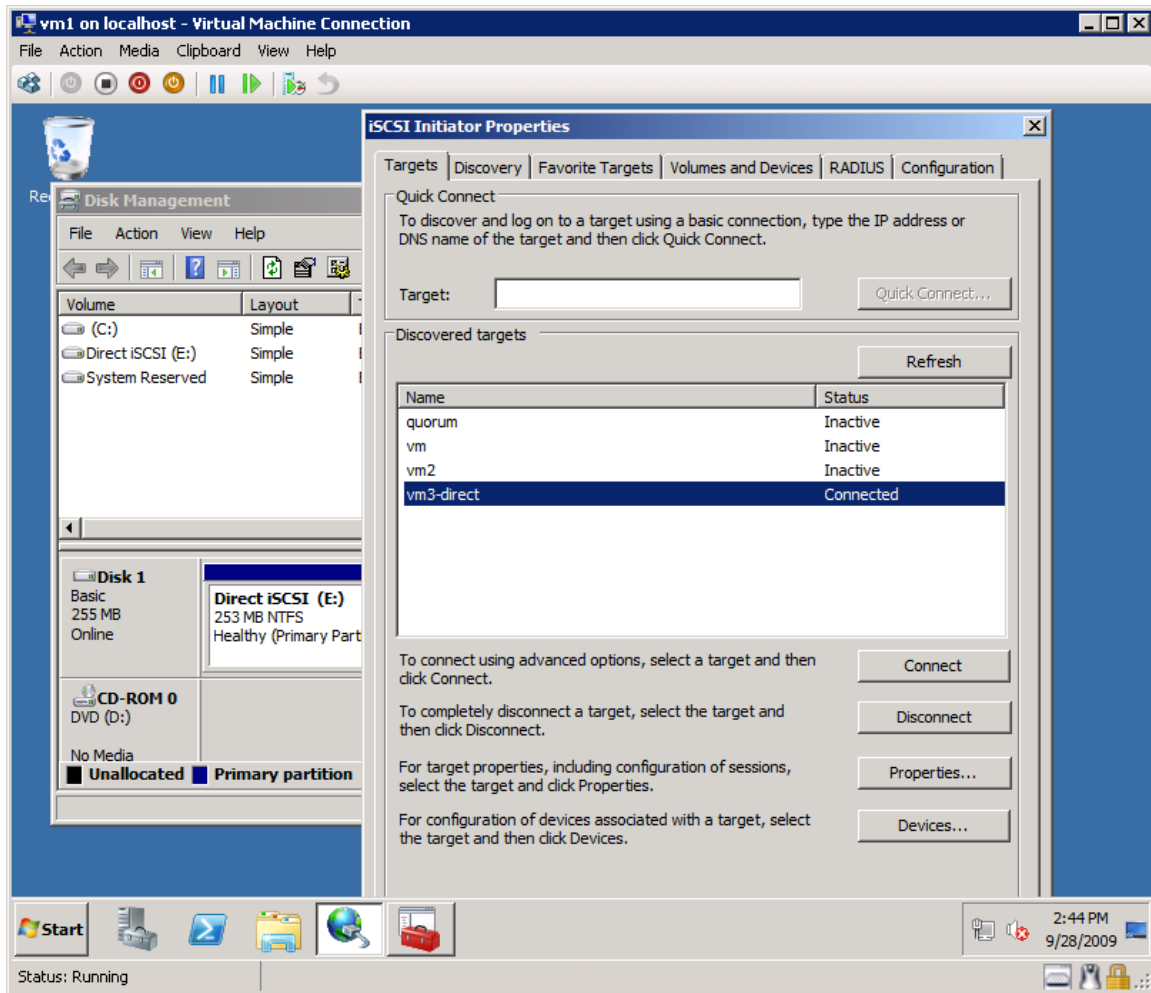
### iSCSI Direct Attachment

Pass-through disks can be an obvious choice when applications require that direct mapping. Yet creating pass-through disks adds a layer of complexity that needn't be present when there aren't specific application requirements. A third option that makes sense for most environments is the direct attachment of iSCSI-based volumes right into the VM. This process uses the VM's iSCSI Initiator to create and manage connections to iSCSI disks.

#### Note

Because direct attachment uses the VM's iSCSI Initiator, this process only works when used with an iSCSI SAN. Environments that use Fibre Channel SANs cannot recognize this benefit and must resort to using pass-through disks.

Figure 1.4 shows how the iSCSI Initiator for VM `\\vm1` is instead configured to connect directly to the previous example's 256MB disk. This connection is possible because of iSCSI's network pervasiveness. Further, the iSCSI Initiator runs as its own service that is independent of the virtualization infrastructure, with the VM's connection to its iSCSI disk being completely isolated from the host.



**Figure 1.4: Connecting directly to an iSCSI LUN from within a VM.**

iSCSI direct attachment enables the highest levels of portability for network-attached disks, retaining all the desired capabilities of the previous examples but without their limitations. Disks can be connected and disconnected at will without the need to reboot the VM. As with VHDs, disks from one VM can be easily attached to an alternate should the need arise; and similar to pass-through disks, data that is contained within the disk remains in its native format.

### **SAN Backups and VM Resources**

When considering the use of a SAN for a virtualized environment, pay special attention to its backup features. One very valuable feature is the ability to directly back up disks without the need for backup agents within the VM. VM-installed agents tend to consume large levels of resources during the backup process, which can have a negative impact on the virtual environment's overall performance. By backing up SAN data directly from the SAN, VMs needn't be impacted by backup operations. This capability represents another benefit to the use of pass-through or direct-attached iSCSI disks.

### **VM-to-VM Clustering**

Yet another capability that can be used by environments with iSCSI SANs is the creation of clusters between VMs. This kind of clustering layers over the top of the clusters used by Hyper-V hosts to ensure the high availability of VMs.

Consider the situation where a critical network resource requires the highest levels of availability in your environment. You may desire that resource to run atop a VM to gain the intrinsic benefits associated with virtualization, but you also want to ensure that the resource itself is clustered to maintain its availability during a VM outage. Even VMs must be rebooted from time to time due to monthly patching operations, so this is not an uncommon requirement. In this case, creating a VM-to-VM cluster for that network resource will provide the needed resiliency.

VM-to-VM clusters require the same kinds of shared storage as do Hyper-V host-to-host clusters. Due to the limitations of the types of storage that can be attached to a VM, that storage can only be created using iSCSI direct connection. Neither VHD attachment nor pass-through disks can provide the necessary shared storage required by the cluster. In this architecture, a SAN disk is exposed and connected to both VMs via direct connection. The result is a network resource that can survive the loss of both a Hyper-V host as well as the loss of a VM.

### **Host Boot from SAN**

With SAN storage becoming so resilient that there is no longer any concern of failure, it becomes possible to move all data off your server's local disks. Eliminating the local disks from servers accomplishes two things: It eliminates the distribution of storage throughout your environment, centralizing everything into a single, manageable SAN solution. Second, it abstracts the servers themselves, enabling a failed server to be quickly replaced by a functioning one. Each server's disk drives are actually part of the SAN, so replacing a server is an exceptionally trivial process.

### **Guest Boot from SAN**

A final solution that can assist with the rapid provisioning of VMs is booting hosted VMs themselves from the SAN. Here, SAN disks are exposed directly to VMs via iSCSI, enabling them to boot directly from the exposed disk. This final configuration is not natively available in Windows Server 2008 R2, and as such requires a third-party solution.

## VM Performance Depends on Storage Performance

As you can see, in all of these architectures, the general trend is towards centralizing storage within the SAN infrastructure. By consolidating your storage into that single location, it is possible to perform some very useful management actions. Storage can be backed up with much less impact on server and VM processing. It can be replicated to alternate or offsite locations for archival or disaster recovery. It can be deduplicated, compressed, thin provisioned, or otherwise deployed with a higher expectation of utilization. In essence, while SAN storage for Hyper-V might be more expensive than local storage, *you should expect to use it more efficiently.*

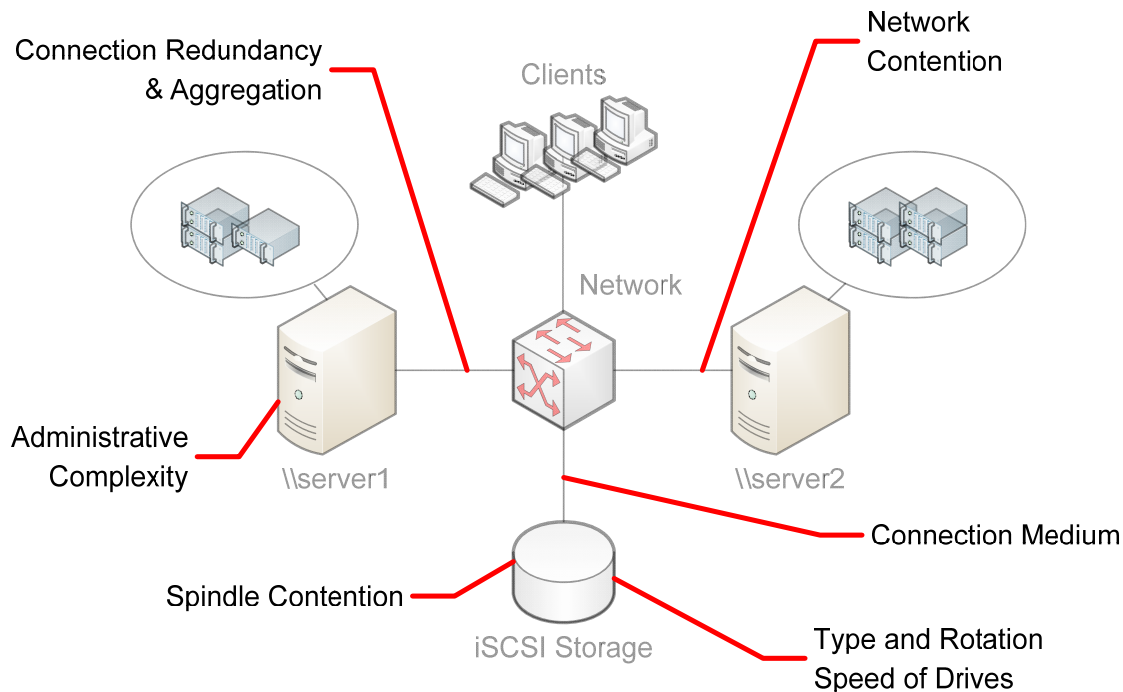
Chapter 3 will focus in greater detail on those specific capabilities to watch for. Yet there is another key factor associated with the centralization of storage that must be discussed here. That factor relates to storage performance.

It has already been said that the introduction of virtualization into an IT environment brings with it added complexities. These complexities arrive due to how virtualization adds layers of abstraction over traditional physical resources. That layer of abstraction is what makes VMs so flexible in their operations: They're portable, they can be rapidly deployed, they're easily restorable, and so on.

Yet that layer of abstraction also masks some of those resources' underlying activities. For example, a virtual network card problem can occur because there is not enough processing power. A reduction in disk performance can be related to network congestion. An entire-system slowdown can be traced back to spindle contention within the storage array. In any of these situations, the effective performance of the virtual environment can be impacted by seemingly unrelated elements.

Figure 1.5 shows how Hyper-V's reliance on multiple, interconnected elements creates multiple points in which bottlenecking can occur. For example, network contention can reduce the amount of bandwidth that is available for passing storage traffic. The type and speed of drives in the storage array can impact their availability. Even the connection medium itself—copper versus fibre, Cat 5 versus Cat 6a—can impact what resources are available to what servers. Smart Hyper-V administrators must always be aware of and compensate for bottlenecks like these in the architecture. Without digging too deep into their technical details, let's take a look at a few that can be common in a Hyper-V architecture.





**Figure 1.5: Virtual environments have multiple areas where performance can bottleneck.**

### Network Contention

Every network connection has a hard limit on the quantity of traffic that can pass along it over a period of time. This maximum throughput is in many environments such a large quantity that monitoring it by individual server is unnecessary. Yet networks that run virtual environments operate much differently than all-physical ones. Consolidating multiple VMs atop a single host means a higher rate of resource utilization (that same “greater efficiency” that was spoken of earlier). Although this brings greater efficiency to those resources, it also brings greater utilization.

Environments that move to virtualization must take into account the potential for network contention as utilization rates go up. This can be alleviated through the addition of new and fully-separated network paths, as well as more powerful networking equipment to handle the load. These paths can be as simple as aggregating multiple server NICs together for failover protection, all the way through completely isolated connections through different network equipment. With Hyper-V’s VMs having a heavy reliance on their storage, distributing the load across multiple paths will become absolutely necessary as the environment scales.

Another resolution involves modifying TCP parameters for specified connections. Microsoft’s Hyper-V R2 supports the use of Jumbo Frames, a modification to TCP that enables larger-sized Ethernet frames to be passed across a network. With a larger quantity of payload data being passed between TCP acknowledgements, the protocol overhead can be reduced by a significant percentage. This results in a performance increase over existing gigabit Ethernet connections.

**Note**

Jumbo Frames are not enabled by default on servers, networking equipment, or most SAN storage devices. Consult your manufacturer's guide for the specific details on how to enable this support. Be aware that Jumbo Frames must be enabled on every interface in each path between servers and storage.

**Connection Redundancy & Aggregation**

Connection redundancy in virtual environments is necessary for two reasons: First, the redundant connection provides an alternate path for data should a failure occur. With external cables connecting servers to storage in iSCSI architectures, the potential for an accidental disconnection is high. For this reason, connection redundancy using MultiPath I/O (MPIO) or Multiple Connected Session (MCS) is strongly suggested. Both protocols are roughly equivalent in terms of effective performance; however, SAN interfaces often support only one of the two options. Support for MPIO is generally more common in today's SAN hardware.

The second reason redundancy is necessary in virtualized environments is for augmenting bandwidth. iSCSI connections to Hyper-V servers can be aggregated using MPIO or MCS for the express purpose of increasing the available throughput between server and storage. In fact, Microsoft's recommendation for iSCSI connections in Hyper-V environments is to aggregate multiple gigabit Ethernet connections in all environments. Ensure that your chosen SAN storage has the capability of handling this kind of link aggregation across multiple interfaces. Chapter 2 will discuss redundancy options in much greater detail.

**Consider 10GbE**

The 10GbE standard was first published in 2002, with its adoption ramping up only today with the increased network needs of virtualized servers. 10GbE interface cards and drivers are available today by most first-party server and storage vendors. Be aware that the use of 10GbE between servers and storage requires a path that is fully 10GbE compliant, including all connecting network equipment, cabling, and interfaces. Cost here can be a concern. With 10GbE being a relatively new technology, its cost is substantially more expensive across the board, with 10GbE bandwidth potentially costing more than aggregating multiple 1GbE connections together.

**Type and Rotation Speed of Drives**

The physical drives in the storage system itself can also be a bottleneck. Multiple drive types exist today for providing disk storage to servers: SCSI, SAS, and SATA are all types of server-quality drives that have at some point been available for SAN storage. Some drives are intended for low-utilization archival storage, while others are optimized for high-speed read and write rates. Virtualization environments require high-performance drives with high I/O rates to ensure good performance of residing VMs.



One primary element of drive performance relates to each drive's rotation speed. Today's SAN drives tend to support speeds of up to 15,000RPM, with higher-rotation speeds resulting in greater performance. Studies have shown that rotation speed of storage drives has a greater impact on overall VM performance than slow network conditions or connection maximums. Ensuring that VHD files are stored on high-performance drives can have a substantial impact on their overall performance.

### Spindle Contention

Another potential bottleneck that can occur within the storage device itself is an oversubscription of disk resources. Remember that files on a disk are linearly written to the individual platters as required by the OS. Hyper-V environments tend to leverage large LUNs with multiple VMs hosted on a single LUN. Multiple hosts have access to those VMs via their iSCSI connections, and process their workloads as necessary.

Spindle contention occurs when too much activity is requested for the files on a small area of disk. For example, if the VHD files for two high-utilization VMs are located near each other on the SAN's disks. When these two VMs have a high rate of change, they require greater than usual attention by the disk's spindle as it traverses the platters to read and write data. When the hardware spindle itself cannot keep up with the load placed upon it, the result is a reduction in storage (and, therefore, VM) performance.

This problem can be easily resolved by re-locating some of the data to another position in the SAN array. Yet, protecting against spindle contention is not an activity that can be easily accomplished by an administrator. There simply aren't the tools at an administrator's disposal for identifying where it is and isn't occurring. Thus, protection against spindle contention is often a task that is automatically handled by the SAN device itself. When considering the purchase of a SAN storage device for Hyper-V, look for those that have the automated capability to monitor and proactively fix spindle contention issues or that use storage virtualization to abstract the physical location of data. Also, work with any SAN vendor to obtain guidance on how best to architect your SAN infrastructure for the lowest-possible level of spindle contention.

### Connection Medium & Administrative Complexity

Lastly is the connection medium itself, with many options available to today's businesses. The discussion on storage in this guide relates specifically to iSCSI-based storage for a number of different reasons: Administrative complexity, cost, in-house experience, and existing infrastructure all factor into the type of SAN that makes most sense for a business:

- **Administrative complexity.** iSCSI storage arrives as a network-based wrapper around traditional SCSI commands. This wrapper means that traditional TCP/IP is used as its mechanism for routing. By consolidating storage traffic under the umbrella of traditional networking, only a single layer of protocols need to be managed by IT operations to support both production and storage networking.

- **In-house experience.** Fibre Channel-based SANs tend to require specialized skills to correctly architect the SAN fabric between servers and storage. These skills are not often available in environments who do not have dedicated SAN administrators on-site. Further, skills in working with traditional copper cabling do not directly translate to those needed for Fibre Channel connections. Thus, additional costs can be required for training.
- **Cost.** Due to iSCSI's reliance on traditional networking devices for its routing, there is no need for additional cables and switching infrastructure to pass storage traffic to configured servers. Existing infrastructure components can be leveraged for all manner of iSCSI traffic. Further, when iSCSI traffic grows to the point where expansion is needed, the incremental costs per server are reduced as well. Table 1.1 shows an example breakdown of costs to connect one server to storage via a single connection. Although cables and Fibre Channel switch ports are both slightly higher with Fibre Channel, a major area of cost relates to the specialized Host Bus Adapter (HBA) that is also required. In the case of iSCSI, existing gigabit copper network cards can be used.

	Fibre Channel	iSCSI
Host Bus Adapter	\$800 to \$2,000 (Fibre Channel HBA)	\$100 to \$200 (GbE Copper NIC)
Cables	\$150	\$15
Switch Port	\$500	\$80
Connection Cost per Server	<b>\$1,450 to \$2,650</b>	<b>\$195 to \$295</b>

**Table 1.1: Comparing the cost to connect one server to storage via Fibre Channel versus iSCSI using a single connection.**

- **Existing infrastructure.** Lastly, every IT environment already has a networking infrastructure in place that runs across traditional copper connections. Along with that infrastructure are usually the extra resources necessary to add connectivity such as cables, switch ports, and so forth. These existing resources can be easily repurposed to pass iSCSI traffic over available connections with a high degree of success.

## iSCSI Makes Sense for Hyper-V Environments

Although useful for environments of all sizes, iSCSI-based storage is particularly suited for those in small and medium enterprises. These enterprises likely do not have the Fibre Channel investment already in place, yet do have the necessary networking equipment and capacity to pass storage traffic with good performance.

Successfully accomplishing that connection with Hyper-V, however, is still more than just a “Next, Next, Finish” activity. Connections must be made with the right level of redundancy as well as the right architecture if your Hyper-V infrastructure is to survive any of a number of potential losses. Chapter 2 will continue this discussion with a look at the various ways to implement iSCSI storage with Hyper-V. That explanation will show you how to easily add the right levels of redundancy and aggregation to ensure success with your Hyper-V VMs.

## Download Additional eBooks from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit <http://nexus.realtimepublishers.com>.