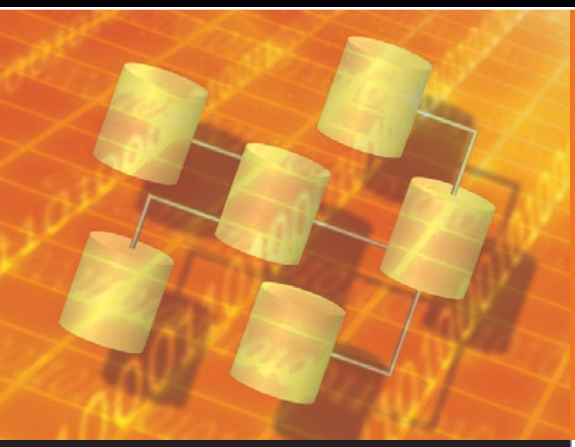

iSCSI: Changing the Economics of Storage

Part 2—Deploying iSCSI in Virtualized Data Centers

BY MATT BAKER

TRAVIS VIGIL

Integrating virtualized servers with shared storage is necessary for flexible virtual machine (VM) mobility, but can be more complicated than many enterprises expect. The second article in this ongoing series details the advantages Internet SCSI (iSCSI) offers in virtualized environments, including simplified deployment, comprehensive storage management and data protection functionality, and seamless VM mobility.

Related Categories:

Dell PowerVault storage

Dell/EMC storage

Internet SCSI (iSCSI)

Storage

Storage architecture

Storage area network (SAN)

Virtualization

Visit www.dell.com/powersolutions for the complete category index.

Virtualization can provide many advantages in enterprise environments, including increased operational efficiency and flexibility, resource utilization, and availability. However, integrating virtualized servers with shared storage to enable flexible virtual machine (VM) mobility can be complex. This is particularly true of traditional Fibre Channel–based storage area networks (SANs), which not only involve specialized knowledge and reconfiguration, but can require sacrificing high-level storage management and data protection functionality to take advantage of virtualization's benefits.

Internet SCSI (iSCSI) can help enterprises avoid these issues. Based on standard Ethernet components, iSCSI-based SANs help simplify configuration and storage management in virtualized environments while avoiding the functionality trade-offs of Fibre Channel. Importantly, iSCSI enables administrators to manage the relationship between VMs and shared storage just as they do with physical servers—a key point for enterprises implementing their first SAN or first virtualized environment and looking for an approach that is familiar, easy to deploy, cost-effective, and able to integrate with their existing storage management and data protection processes.

Part 1 of this series[1] outlined the basics of iSCSI, which allows enterprises to create Ethernet-based SANs at a lower cost of entry than Fibre Channel without requiring specialized equipment or expertise, and included a discussion of how enterprises can implement iSCSI using Dell™ PowerVault™ and Dell/EMC storage arrays. This second installment focuses on how administrators can combine iSCSI and virtualization to build an efficient data center that is optimized for enterprise workloads.

Understanding virtualization and storage

Virtualization software enables multiple isolated environments to run in individual VMs on the same physical server, with each VM having its own virtual BIOS, processors, memory, and so on. Abstracting the OS platform from the hardware allows administrators to easily consolidate physical resources and flexibly move VMs between servers as needed for load balancing and failover.

Along with these advantages, however, virtualization can bring its own set of challenges—and integrating VMs with a shared storage back end can be a cause of administrator headaches. Fibre Channel–based SANs can be complex, particularly for enterprises deploying shared storage for the first time. Fibre Channel does offer many advantages, in particular high performance for high-throughput, low-latency

[1] "iSCSI: Changing the Economics of Storage; Part 1—Understanding iSCSI in Enterprise Environments," by Travis Vigil, in *Dell Power Solutions*, May 2007, www.dell.com/downloads/global/power/ps2q07-20070335-Vigil.pdf.

applications such as large-scale decision support. However, iSCSI can help enterprises address three common issues encountered in virtualization deployments:

- **Managing complex storage relationships:** SANs can be complex to deploy and manage on their own. By creating a large number of VM-to-storage relationships that must be managed by the hypervisor—the software layer providing the virtualized environment—virtualization can increase that complexity.
- **Enabling storage management and data protection:** Storage management and data protection operations must be re-created in the hypervisor when using Fibre Channel–based shared storage. Not only does this configuration require specialized knowledge, but in some cases, array-based functionality may not be available with the hypervisor, and scripts developed for some backup applications in non-virtualized environments may no longer function or may need to be rewritten entirely in virtualized environments.
- **Enabling VM mobility:** The ability to move VMs freely between physical servers is essential for load-balancing and failover capabilities. However, enabling this mobility on Fibre Channel–based SANs can result in administrative complexity for enterprises unfamiliar with managing shared storage for virtualized environments.

> "iSCSI-based SANs help simplify configuration and storage management in virtualized environments while avoiding the functionality trade-offs of Fibre Channel."

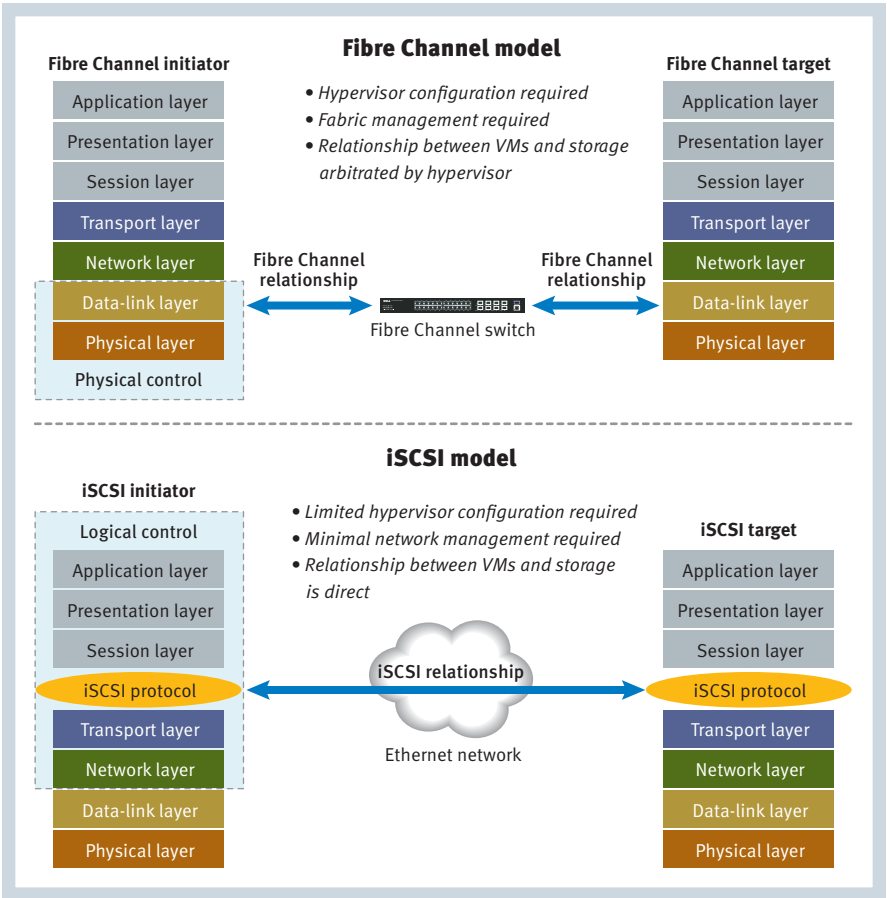

Figure 1. *Fibre Channel and iSCSI storage management models*

In the future, initiatives such as N_Port ID Virtualization (NPIV)—which allows multiple Fibre Channel initiators to share a single physical port—may help simplify the configuration and management of Fibre Channel–based SANs, particularly in virtualized environments. However, iSCSI can mitigate or help eliminate these problems now, allowing administrators to benefit from the advantages of SANs and virtualization while helping simplify deployment and management—and avoiding the functionality limitations currently incurred in a Fibre Channel–based virtualization deployment.

Comparing Fibre Channel and iSCSI in virtualized environments

iSCSI can help enterprises meet all three of the major challenges listed in the preceding section, streamlining the management of complex storage relationships, allowing the same comprehensive storage management and data protection in both virtualized and non-virtualized environments, and enabling easy VM mobility with less storage-specific hypervisor configuration than can be required with Fibre Channel.

Managing complex storage relationships

One difference between Fibre Channel and iSCSI in virtualized environments lies in how they construct the relationships between VMs and storage. Although neither strictly adheres to the Open System Interconnection (OSI) model, it can be helpful to think of Fibre Channel as controlled at the lowest OSI layers (the physical and data-link layers) and iSCSI as controlled at the higher, logical layers, as shown in Figure 1.

The difficulties of managing Fibre Channel–based storage in virtualized environments stem primarily from its *hypervisor-managed storage* approach, in which the relationships between VMs and storage are arbitrated by the hypervisor (see Figure 2). Hypervisor-managed storage does offer the advantage of quick and easy VM creation, and can be particularly useful when each VM has only one virtual disk, a common setup for non-data-intensive workloads. Hypervisor-managed storage also offers the advantage of hypervisor-specific backup and data-protection capabilities (such as those of VMware® Consolidated Backup), which are optimized for quick and easy mass backups of VMs' storage. But hypervisor-managed storage can also require significant configuration, especially when the VMs are running applications that must access large amounts of data. For example, administrators must first assign storage resources to the World Wide Name (WWN) associated with the physical server, then use the virtualization software to provision these resources to the VMs.

In contrast, software initiators allow iSCSI to function independently of the hypervisor. Hypervisor-managed storage is still required for boot logical units (LUNs), and can offer iSCSI the same benefits it does for Fibre Channel. However, using iSCSI also allows administrators to take advantage of the *storage direct* approach (see Figure 3). This model enables one of the most compelling benefits of iSCSI in virtualized environments: creating direct relationships between VMs and storage without hypervisor configuration. In doing so, iSCSI streamlines network management, allowing administrators to treat the storage associated with VMs the same way they treat storage in non-virtualized environments and dramatically simplifying SAN configuration and management. For data-intensive workloads such as Microsoft® Exchange and SQL Server™ software, this simplification can be crucial.

Figure 4 summarizes some of the features and advantages of two examples of hypervisor-managed storage—VMware Virtual Machine File System (VMFS) and raw device mapping (RDM)—and the storage direct method.

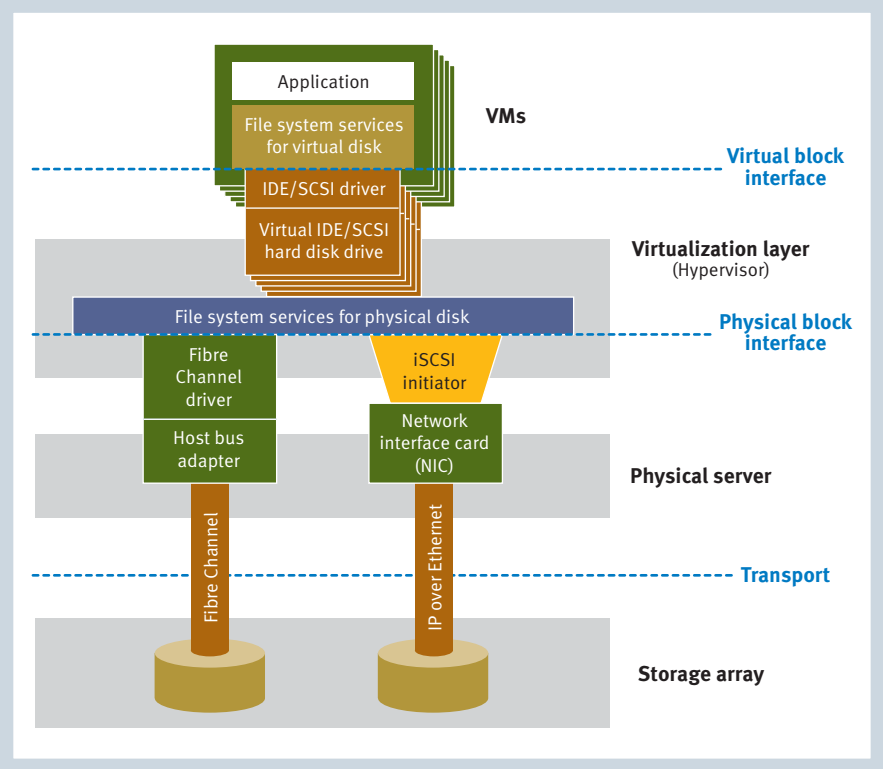

Figure 2. *Hypervisor-managed storage architecture in a virtualized environment*

Enabling storage management and data protection

Storage management and data protection tools—such as array management, snapshot, and backup and recovery software—are essential to many enterprises, allowing administrators to efficiently monitor and manage storage resources and helping ensure the availability of critical data. But use of this functionality may be limited in certain ways when using Fibre Channel:

- **Unfamiliar methods:** Administrators cannot use these tools in virtualized environments the same way they do in non-virtualized environments. As a result, they must learn and implement different ways to manage, back up, and recover storage.
- **Limited functionality:** Even after carrying out the reconfiguration necessary to re-create standard functionality at the hypervisor level, administrators may discover that some functionality readily available in non-virtualized environments is still not available after the move to a virtualized environment. In particular, array-based management, snapshot integration into popular software like Microsoft Exchange and SQL Server, and multipath functionality can be limited. In addition, backup applications or scripts that utilize Microsoft Volume Shadow Copy Service (VSS) do not function when run in the hypervisor unless the independent software vendors have enabled them to do so and support VSS in the hypervisor.
- **Difficult migrations:** Migrating applications, images, and scripts from a physical server to a VM, from a VM to a physical server, or between two different virtualized environments over Fibre Channel can require reconfiguration and re-qualification by expert administrators, and any errors may cause migrated applications to fail.

The iSCSI storage direct approach is designed to significantly simplify the use of storage management and data protection tools in virtualized environments:

- **Familiar methods:** Administrators can easily use storage management software and agents—such as VSS providers and the EMC® Navisphere® command-line interface—and create direct backups to tape or disk the same way they would in non-virtualized environments.
- **Robust functionality:** The storage direct method helps preserve array-based management functionality, including single-mailbox restores for Microsoft Exchange servers, single-LUN restores, point-in-time snapshots, and so on. Comprehensive multipath and scripting capabilities are designed to function the same way in both virtualized and non-virtualized environments.
- **Streamlined migrations:** Applications, images, and scripts developed on a physical server can function on a VM the same way they do on the physical server without complex reconfiguration, and vice versa.

Figure 5 summarizes characteristic differences between Fibre Channel and iSCSI for storage management and data protection tools in virtualized environments.

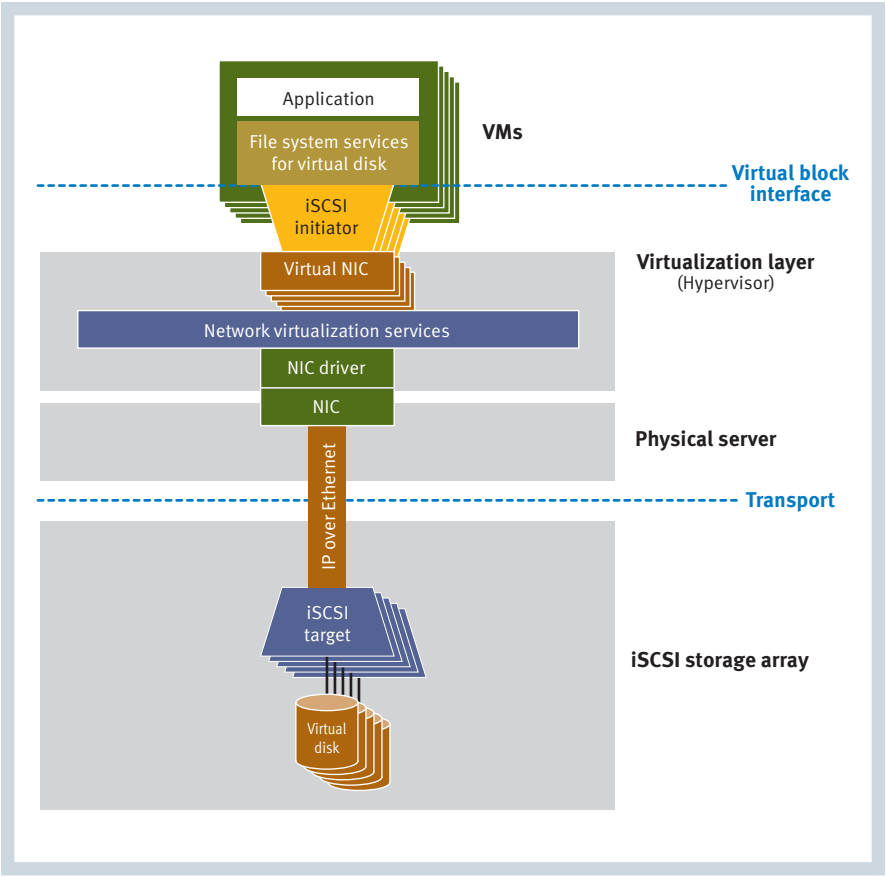

Figure 3. *Storage direct architecture in a virtualized environment*

Enabling virtual machine mobility

The ability to move VMs across physical servers quickly and easily to accommodate load balancing and failover is a key advantage of virtualization—provided the infrastructure can accommodate dynamic resource allocation. As with the relationship between VMs and storage, however, enabling VM mobility in Fibre Channel–based SANs typically requires hypervisor arbitration to allow VMs to maintain storage access from different physical servers. In addition, enabling VM mobility can raise concerns about data integrity due to administrator mis-configuration: giving every server access to all storage resources requires administrators to have multiple servers accessing the same LUN.

Because each iSCSI initiator and target have a unique name—usually the iSCSI Qualified Name (IQN)—each VM has a unique, direct relationship with storage resources

	Hypervisor-managed storage		Storage direct
	VMFS	RDM	
Support	Fibre Channel and iSCSI	Fibre Channel and iSCSI	iSCSI only
Advantages	• Allows easy VM creation • Provides basic snapshot functionality, and can be used with VMware Consolidated Backup for quick and easy mass backup of VMs • Requires minimal administrator interaction with guest operating systems • Enables access to array-based applications (with limitations)	• Allows 1:1 LUN mapping similar to storage direct • Provides more granular snapshot and clone functionality than VMFS • Requires minimal administrator interaction with guest operating systems • Enables greater access to array-based applications than VMFS (but still with some limitations) • Offers increased performance over VMFS	• Provides the most granular snapshot and clone functionality of these methods • Can integrate snapshots and clones with Microsoft applications based on Microsoft VSS and Virtual Disk Service (VDS) backups—such as Exchange and SQL Server—allowing administrators to use familiar tools • Enables full access to array-based applications, and supports robust multipath functionality for load balancing and failover • Offers increased performance compared with hypervisor-managed iSCSI • Supports secure connections between guest operating systems and targets as well as seamless VM mobility • Simplifies virtual-to-physical and physical-to-virtual migrations by helping eliminate the need to retest and reconfigure storage

Figure 4. *Comparison of the VMware VMFS and RDM hypervisor-managed storage methods and the storage direct method*

	Fibre Channel	iSCSI (storage direct)
Storage management	Managed at both the VM and hypervisor level, with limited functionality	Directly managed at the VM level, with comprehensive functionality
Backup and recovery	Managed and arbitrated at the hypervisor level	Enabled through a direct relationship between the VM and the tape or disk storage
Multipath	Functionality typically limited to failover	Comprehensive functionality available
Migration	Applications, images, and scripts developed on physical servers can require significant modification to run on VMs, and vice versa	Applications, images, and scripts developed on physical servers can run seamlessly on VMs, and vice versa

Figure 5. *Comparison of storage management and data protection capabilities in virtualized environments using Fibre Channel and iSCSI*

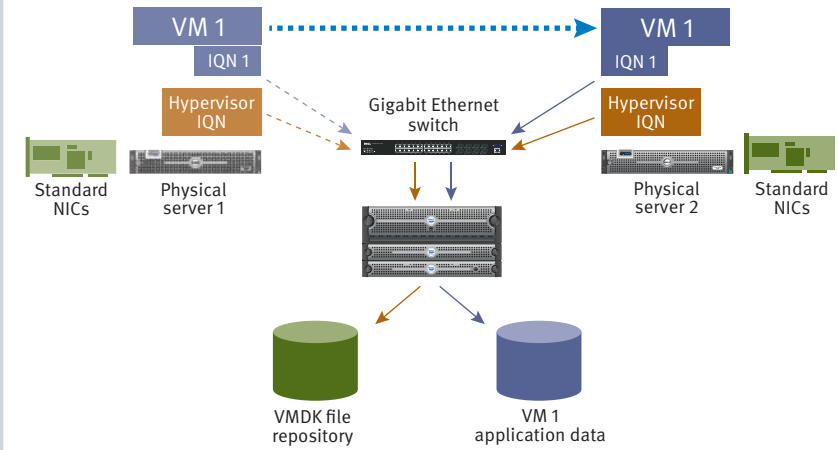

Figure 6. *Seamless virtual machine mobility using the iSCSI storage direct approach*

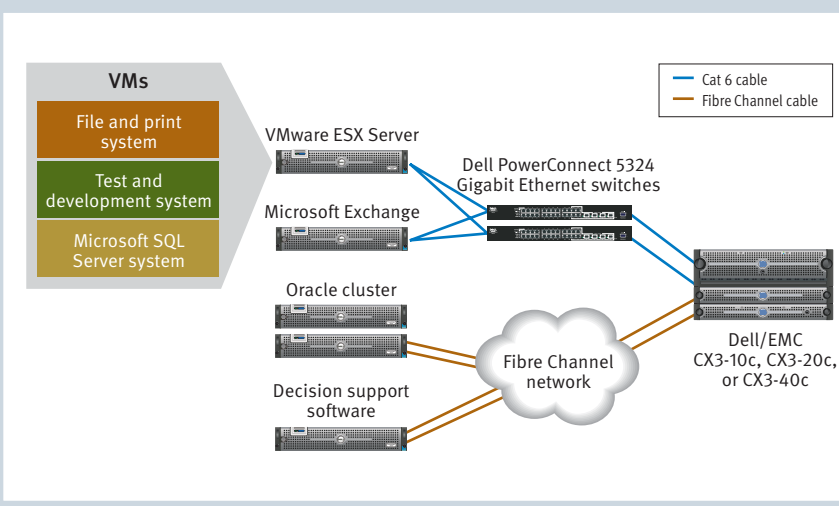

Figure 7. *Example environment combining Fibre Channel and iSCSI with Dell/EMC storage in a tiered storage structure*

distinct from the relationship between the physical server and storage, helping preserve tight access control. And because this relationship persists independently of the physical host and the hypervisor, VMs can move seamlessly between physical servers with less storage-specific hypervisor configuration than can be required with Fibre Channel. Figure 6 illustrates this process: to balance workloads or move them for proactive maintenance, a VM can easily migrate to another server while maintaining its unique identity and storage access on the new host.

Building a tiered data center

Enterprises can take advantage of storage arrays offering iSCSI functionality, such as the Dell/EMC AX150i, Dell/EMC CX3 series, and Dell PowerVault NX1950, to create a tiered storage architecture designed to optimize resources based on server workload. For example, an investment in Fibre Channel is typically appropriate for high-throughput, low-latency applications running on non-virtualized servers, while iSCSI offers distinct advantages for applications with random I/O and those running on virtualized servers. To provide additional flexibility for enterprises creating such a tiered environment, the Dell/EMC CX3-10c, CX3-20c, and CX3-40c offer both Fibre Channel and iSCSI connectivity.

Figure 7 illustrates an example optimized environment deploying both Fibre Channel and iSCSI. The file and print, test and development, and Microsoft SQL Server systems run on VMs based on the VMware ESX Server virtualization platform, and use iSCSI to take advantage of its strengths in virtualized environments. Similarly, the Microsoft Exchange server uses iSCSI to take advantage of its cost-effectiveness and performance for random I/O workloads. The Oracle® database cluster and decision support software, meanwhile, run directly on physical servers and utilize Fibre Channel to take advantage of its strengths for those types of applications.

Figure 8 illustrates another type of tiered environment using the flexible Dell PowerVault NX1950 networked storage solution, which can store both file and application data and is

designed to work with multiple operating environments and communication protocols. In the example environment, the virtualized file and print, test and development, and Exchange systems use iSCSI for block-level access to the storage. The clients running the Microsoft Windows® and Linux® operating systems, meanwhile, use Common Internet File System (CIFS) and Network File System (NFS) for file-level access and sharing. Administrators can use the central console in the included Microsoft Windows Unified Data Storage Server 2003 OS to easily create and manage file shares, iSCSI targets, point-in-time snapshots, performance logs and metrics, and more.

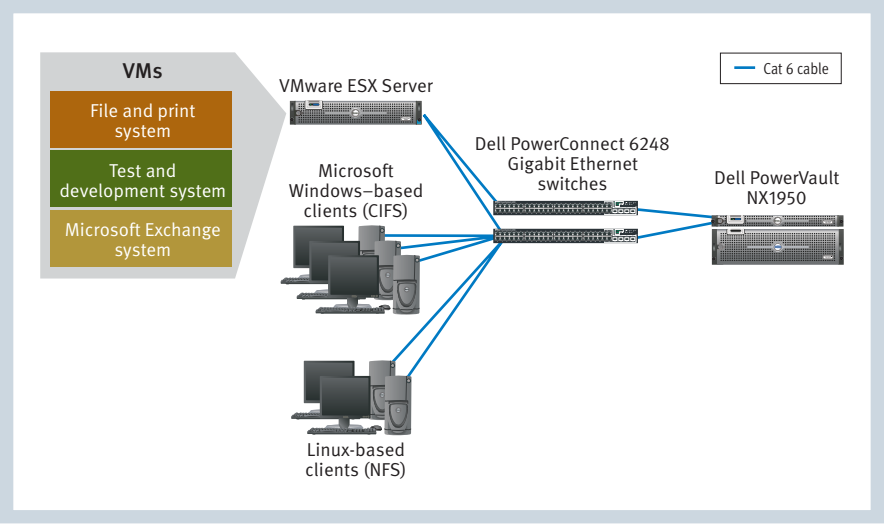

Figure 8. *Example environment combining block-level access over iSCSI for servers and file-level access for clients with the Dell PowerVault NX1950*

Deploying flexible, cost-effective virtualization

Enterprises can face multiple challenges when integrating a virtualized environment with a shared storage back end. Fibre Channel may be the best choice for virtualization for enterprises that have already broadly deployed Fibre Channel–based SANs or those whose environments require high performance. However, combining Fibre Channel with virtualization may require specialized expertise to re-create basic functionality at the hypervisor level and reconfiguration to enable VM mobility for load balancing and failover. In addition, using Fibre Channel with virtualization may limit array-based management, backup and recovery, and multipath functionality.

Besides allowing enterprises to create SANs using standard, familiar Ethernet equipment with a lower cost of entry than Fibre Channel, iSCSI can offer significant advantages in virtualized environments. Its potential benefits include simplified deployment based on direct relationships between VMs and storage, which allows administrators to manage VMs the same way they do physical servers; comprehensive storage management and data protection for both VMs and physical servers; and seamless VM mobility without complex reconfiguration. By deploying iSCSI on its own or in addition to Fibre Channel, enterprises can easily integrate virtualized servers with shared storage and create a tiered environment optimized for their particular needs.

Matt Baker *is a product marketing strategist in the Dell Storage Product Group, where he analyzes emerging storage technologies and market trends and conducts long-range product planning. Before joining Dell, Matt spent 10 years with Intel in various IT, product engineering, and marketing roles, focusing mainly on traditional IP networking and IP-based storage networking.*

Travis Vigil *is a senior product marketing consultant for Dell iSCSI and Dell PowerConnect™ solutions. He has nearly 10 years of experience with technology companies including Intel and Dell, and was most recently the product manager for Dell PowerVault disk storage. He has a B.S. from Stanford University and an M.B.A. from Northwestern University's Kellogg School of Management.*

more Online

www.dell.com/**powersolutions**

QUICK LINK

Dell iSCSI storage solutions:
www.dell.com/iscsi