WHITE PAPER

VIRTUALIZED iSCSI SANs: Flexible, Scalable Enterprise Storage for Virtual Infrastructures

vmware®

DELL EQUALLOGIC®

MARCH 2008

Table of Contents

Server Virtualization 3
Virtual Infrastructure 4
Storage Virtualization 5
Requirements for Networked Enterprise Storage in a VMware Infrastructure 5
What is iSCSI? 6
ESX Server and iSCSI 7
How iSCSI Enables Virtualized SANs 7
EqualLogic Virtualized iSCSI SANs 8
Automatic Load Balancing 9
Non-Disruptive Scalability 9
Quick, Intellegent Provisioning 9
Automated Management 10
A flexible, Scalable, Virtualized Environment for Servers and Storage 10
Summary 11
Customer Case Study 12

VIRTUALIZED iSCSI SANS: Flexible, Scalable Enterprise Storage for Virtual Infrastructures

iSCSI is a flexible and powerful storage area networking (SAN) protocol that delivers superior capabilities and benefits across all market segments. In addition to providing enterprise-class data availability and performance, the iSCSI protocol enables breakthrough virtual storage designs that parallel the advanced designs of server virtualization technologies such as VMware® Infrastructure 3.

Enterprises of all sizes are building flexible storage infrastructures using iSCSI and advanced virtualization technologies that let them allocate and shift SAN resources dynamically to respond to the demands of their virtualized server environments.

This white paper describes a virtualized infrastructure that applies storage and server virtualization technologies to cost-effectively achieve a flexible, high-performance, dynamic IT infrastructure that is simple to manage and scale.

Server Virtualization

Virtualization is a proven software technology that is rapidly transforming the IT landscape and fundamentally changing the way that people compute.

Today's powerful x86 computer hardware was originally designed to run only a single operating system and a single application. Virtualization breaks that bond, making it possible to run multiple operating systems and multiple applications on the same computer at the same time, increasing the utilization and flexibility of hardware.

In essence, virtualization lets you transform hardware into software. You can use software such as VMware® ESX Server to transform or "virtualize" the hardware resources of an x86-based computer—including the CPU, RAM, hard disk and network controller—to create a fully functional virtual machine that runs its own operating system and applications, just like a "real" computer.

In general, VMware virtual machines possess four key characteristics that benefit the user:

- **Compatibility:** Just like a physical computer, a virtual machine hosts its own guest operating system and applications, and has all the components found in a physical computer (motherboard, VGA card, network card controller, etc). As a result, virtual machines are completely compatible with all standard x86 operating systems, applications and device drivers, so you can use a virtual machine to run all the same software that you would run on a physical x86 computer.

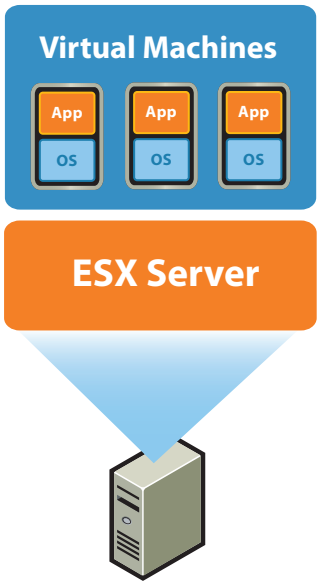

Physical Server

- **Isolation:** While virtual machines can share the physical resources of a single computer, they remain completely isolated from each other as if they were separate physical machines. If, for example, there are four virtual machines on a single physical server and one of the virtual machines crashes, the other three virtual machines remain available. Isolation is an important reason why the availability and security of applications running in a virtual environment is far superior to applications running in a traditional, non-virtualized system.
- **Encapsulation:** A virtual machine is essentially a software container that bundles or "encapsulates" a complete set of virtual hardware resources, as well as an operating system and all its applications, inside a software package. Encapsulation makes virtual machines incredibly portable and easy to manage. For example, you can move and copy a virtual machine from one location to another just like any other software file, or save a virtual machine on any standard data storage medium, from a pocket-sized USB flash memory card to enterprise storage area networks (SANs).

- Hardware independence: Virtual machines are completely independent from their underlying physical hardware. For example, you can configure a virtual machine with virtual components (e.g., CPU, network card, SCSI controller) completely different than the physical components that are present on the underlying hardware. Virtual machines on the same physical server can even run different kinds of operating systems (Windows, Linux, etc).

  When coupled with the properties of encapsulation and compatibility, hardware independence gives you the freedom to move a virtual machine from one type of x86 computer to another without making any changes to the device drivers, operating system or applications. Hardware independence also means that you can run a heterogeneous mixture of operating systems and applications on a single physical computer.

Multiple virtual machines share hardware resources without interfering with each other so that you can safely run several operating systems and applications at the same time on a single computer.

The VMware approach to virtualization inserts a thin layer of software directly on the computer hardware or on a host operating system. This software layer creates virtual machines and contains a virtual machine monitor or "hypervisor" that allocates hardware resources dynamically and transparently so that multiple operating systems can run concurrently on a single physical computer without even knowing it.

However, virtualizing a single physical computer is just the beginning. VMware offers a robust virtualization platform that can scale across hundreds of interconnected physical computers and storage devices to form an entire virtual infrastructure.

Virtual Infrastructure

The introduction of virtualization technology presents a number of opportunities for driving capital and operational efficiency above and beyond the simple benefit of safe partitioning. VMware customers have harnessed the power of virtualization to better manage IT resources, provide better service levels and streamline IT processes.

In essence, a virtual infrastructure is a dynamic mapping of physical resources to business needs. While a virtual machine represents the physical resources of a single computer, a virtual infrastructure represents the physical resources of the entire IT environment, aggregating x86 computers and their attached network and storage into a unified pool of IT resources.

Structurally, a virtual infrastructure consists of the following components:

- A virtualization hypervisor, configured on each X86-based server, fully utilizing the underlying hardware infrastructure.
- A set of virtualization-based distributed system infrastructure services, such as resource management, to optimize available resources among virtual machines.
- Automation solutions that provide special capabilities to optimize a particular IT process, such as provisioning or disaster recovery.

By decoupling the entire software environment from its underlying hardware infrastructure, virtualization allows for the aggregation of multiple servers, storage infrastructure and networks into shared pools of resources that can be delivered dynamically, securely and reliably to applications, as needed. This pioneering approach enables organizations to build a computing infrastructure with high levels of utilization, availability, automation and flexibility using building blocks of inexpensive, industry-standard servers.

VMware Infrastructure

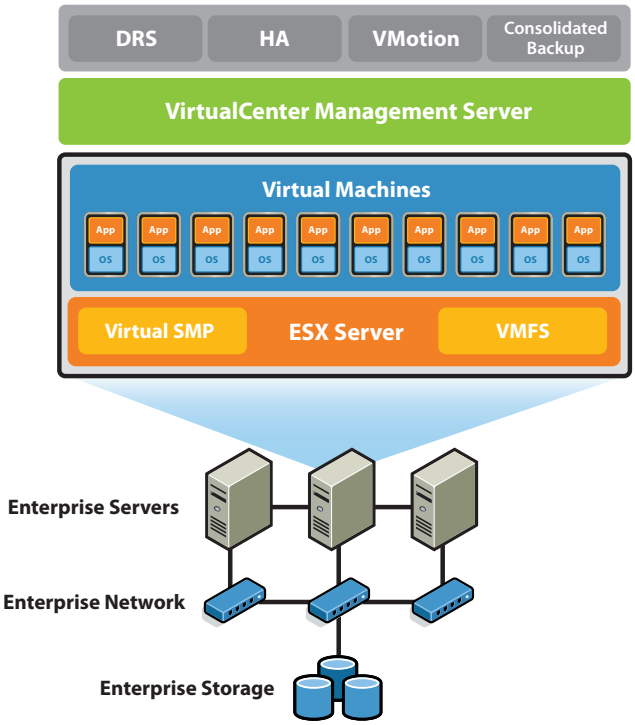

VMware has made it possible to fully realize the enormous benefits of virtualization in production-scale IT environments by building virtual infrastructure automation and management capabilities around a best-in-class hypervisor. In fact, 86 percent of VMware customers use virtualization in production, and 43 percent deploy most new production applications in virtual machines.

VMware virtual infrastructure solutions are ideal for production environments, in part because they run on industry-standard servers and desktops and support a wide range of operating system and application environments as well as networking and storage infrastructures. VMware designs its solutions to function independently of the hardware and operating system to provide customers with a broad platform choice. As a result, VMware solutions provide a key integration point for hardware and infrastructure management vendors to deliver differentiated value that can be applied uniformly across all application and operating system environments.

VMware customers who have adopted virtual infrastructure solutions have reported dramatic results, including:

- 60-80% utilization rates for x86 servers (up from 5-15% in non-virtualized PCs).
- Cost savings of more than $3,000 annually for every workload virtualized.
- The ability to provision new applications in minutes instead of days or weeks.
- An 85 percent improvement in recovery time from unplanned downtime.

Storage Virtualization

Today, most IT professionals define storage virtualization as a technology that allows discrete storage systems to operate as a single resource. In light of recent advancements in server virtualization, the concept of storage virtualization is being further refined as a way to create an abstraction layer between the storage hardware and logical data volumes.

Given sufficient protocol support (as with iSCSI), virtual storage products are now being designed that allow data volumes to be located and striped across multiple (and diverse) physical storage resources, including storage systems, RAID groups, disk types and controllers. Not only does this provide higher performance and scalability, it also allows data volumes to be transparently moved from one set of resources to another without disruption to the operating systems and applications that are using data. It forms a scalable, highly resilient, flexible storage environment, resulting in better storage utilization rates and far lower operating costs.

Requirements for Networked Storage in a VMware Infrastructure

In a VMware Infrastructure 3 environment, multiple physical servers are networked into a cooperating set of computing resources (see Figure 1). In lieu of a standard Linux or Windows® operating system, each of the x86 standard servers run a copy of VMware ESX Server, a robust, production-proven virtualization layer that abstracts processor, memory, storage and networking resources into multiple virtual machines. The VMware Infrastructure suite allows many x86 physical servers to be combined into a single resource pool that aggregates processor, memory, disk and networking capacity. Virtual machines are deployed to the resource pool rather than to particular machines. VMware® VirtualCenter delivers centralized management, operational automation, resource optimization and high availability to IT environments based on VMware Infrastructure.

Within the VMware environment, a virtual machine's configuration parameters, operating system, data sets and applications are encapsulated in a set of files managed by the VMFS file system. VMFS is a high-performance clustered file system that allows multiple ESX Servers to access the same virtual machine storage concurrently. It enables virtualization-based distributed services such as live migration of running virtual machines, automatic restart of a failed virtual machine on a different physical server, and clustering of virtual machines across different physical servers.

The power of this architecture comes from the ability to run any guest OS, its applications and its data without modifying any physical server in the cluster. As the need arises, additional server resources can be non-disruptively added to the ESX Server cluster, and current workloads are load balanced to take advantage of the newly available resources.

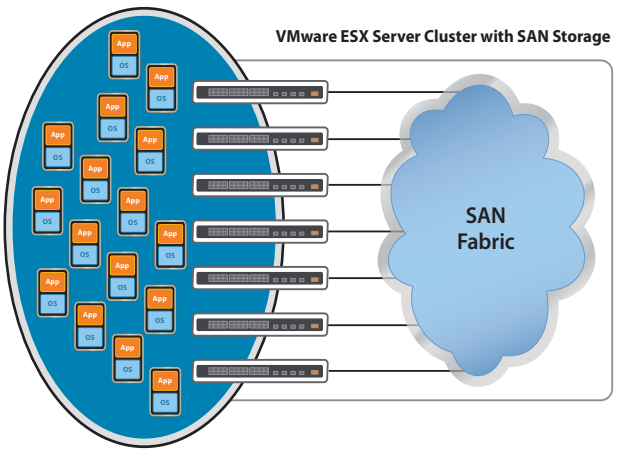

WHAT IS iSCSI?

Developed by the Internet Engineering Task Force (IETF), iSCSI is a standard storage access protocol for interconnecting servers and storage using an IP-based network interconnect. iSCSI is built upon the SCSI and TCP/IP protocols, the dominant standards for storage and networking in use today. iSCSI is recognized for its simple integration, interoperability, fundamentally lower costs and its ability to leverage an organization's inherent networking skills for quick and broad adoption.

Software-based initiators, generally available for all major operating systems at minimal or no cost, take advantage of standard Ethernet interfaces included as basic features of servers and desktops, enabling widespread deployment. Easy access to this technology at little to no cost, coupled with the application of pre-existing IP networking skills fundamental to today's IT organizations, provides for easy, affordable and rapid adoption of iSCSI technology.

Using an ordinary IP network, iSCSI transports block-level data between an iSCSI initiator on a server and an iSCSI target on a storage device. The iSCSI protocol encapsulates SCSI commands and assembles the data into TCP/IP packets sent over the network using a point-to-point connection. Upon arrival at the storage device, the encapsulated SCSI commands are unpacked from the TCP/IP packet for processing.

When an iSCSI initiator connects to an iSCSI target, the storage is seen by the operating system as a local SCSI device that can be formatted as usual. The process is transparent to applications, file systems and operating systems. By consolidating storage with an iSCSI Storage Area Network (SAN), multiple platforms can share the same storage, greatly improving utilization and efficiency.

The opportunity to use a standard Gigabit Ethernet NIC to connect servers to storage makes iSCSI both simple and affordable. Today's high-speed CPUs can run iSCSI at line speed over standard NICs using only a marginal amount of CPU. Onboard Gigabit Ethernet NICs ship standard with servers, so no additional cost is incurred and the administrator need not bother with installing adaptor cards.

Together, standard NICs and the lower cost of Ethernet switches offer a compelling cost savings over Fibre Channel, while maintaining the performance advantages of a SAN. However, the savings are not limited to the hardware costs alone. With iSCSI, IT departments leverage their existing IP networking expertise, without needing specially trained staff. This familiarity with the network infrastructure eliminates a key problem of complexity associated with Fibre Channel SANs.

Effective VMware deployments depend on shared access to storage – in other words, a SAN. SANs ensure that each ESX Server has immediate access to any virtual machine's data sets, enabling immediate re-hosting of the virtual machine. This eliminates the unacceptable and time consuming necessity of copying virtual machine files, applications, and data from one ESX Server to another.

SANs support powerful features within VMware Infrastructure – these include VMware® VMotion, the online migration of active virtual machines without interruption, and VMware® Distributed Resource Scheduler (DRS), which continuously balances virtual machine workloads across resource pools based on rules that can be adjusted as conditions change. When load increases, VMware DRS automatically allocates additional resources and uses VMware VMotion to transparently migrate virtual machines between hosts in the resource pool to ensure that service level agreements are met. VMware® High Availability (HA) monitors the proper functioning of ESX Servers within the cluster, and in case of an server failure, re-hosts and restarts virtual machines affected by the failure onto another available server within the cluster. VMware HA provides uniform high availability across the entire virtualized IT environment without the cost and complexity of failover solutions tied to either operating systems or specific applications.

In addition to the requirement for networked storage, a virtualized environment heightens the need for high-performance, highly available, resilient storage to meet the needs of the aggregation of workloads. As more critical applications, production workloads, and data assets are consolidated into fewer resources, the need for high-performance, non-disruptive scalability, and continuous availability of the storage assets increases. Enterprise-class storage, designed for mission-critical deployments, is a basic requirement when building a virtualized IT infrastructure for production use.

In particular, organizations should consider deploying purpose-built storage architectures that include mirrored memory write caches, fully-redundant hot-pluggable components, online hot spare disks, environmental monitoring and enterprise-class disk drives with RAID protection. Furthermore, advanced availability features, including storage controller and I/O path failover, are recommended to guarantee data access, even in the case of component failure.

Companies should also evaluate the data protection software features of the storage system. Storage consolidation within a SAN allows for the consistent application of data protection and disaster recovery, assuming the basic SAN-based tools are available and enabled. Space-efficient, non-intrusive, snapshot based point-in-time copies, as well as efficient array-based replication tools, should be considered as basic requirements of the storage infrastructure for a virtualized environment. Multiple use cases exist for the application of these tools in virtualized environments.

- Disaster recovery
- Online backup and quick recovery of virtual machines
- Simple extraction of lost or corrupted data files
- Rapid virtual machine and data set provisioning using snapshots and clones for production and temporary use
- Server-less tape- and disk-based backup and recovery of the organization's data assets

A thorough evaluation of the storage platform should be considered prior to deployment, reflecting not only upon the robustness and feature set, but also the total cost of ownership over its useful life, with particular consideration of the required system growth in terms of both performance and capacity.

ESX Server and iSCSI

With the advent of VMware Infrastructure 3, iSCSI initiators are integrated into the ESX kernel, allowing native access to iSCSI storage directly from ESX Servers. Both hardware initiators (QLogic™ QLA404xC, QLA406xC) and software initiators are available. NIC teaming (multiplexing a single logical connection across multiple interfaces) is supported with the software initiator. The hardware initiator includes support for multi-pathing as well as boot from SAN.

iSCSI support broadens the potential for full-fledged VMware Infrastructure 3 deployments in small and medium businesses by obviating the need for a costly and complex Fibre Channel SAN deployment. The acquisition, implementation and operational cost of an iSCSI-based SAN are intrinsically lower, removing an economic barrier to the adoption of server virtualization technologies. In addition, but far less well recognized, is the opportunity for advanced iSCSI-based system architectures to dramatically change customer expectations in mid-market and enterprise deployments, as well.

Adding iSCSI protocol interfaces to classic storage system designs does not fully exploit the potential of the iSCSI protocol and, in turn, does not adequately solve the storage management challenges faced today by IT managers today. Adding iSCSI protocol interfaces to classic storage system designs does not fully exploit the potential of the iSCSI protocol and, in turn, does not adequately solve the storage management challenges faced by IT managers today. Storage administrators require a rarified level of knowledge for configuring and tuning storage arrays, RAID geometries and data layout considerations, application workload analysis, forced data migration and complex system upgrades.

How iSCSI Enables Virtualized SANs

Storage virtualization can simplify provisioning and ongoing management, increase storage utilization, provide unlimited scalability in capacity and performance, and enable online migration of data sets among controllers and storage tiers. Network-based storage is consolidated into a simple, flexible pool of storage which grows in capacity and performance (for example, in a Dell EqualLogic environment). Storage virtualization then becomes a key enabler for simplifying a virtualized IT infrastructure.

Historically, the adoption of storage virtualization technologies in SAN environments has been limited primarily to enterprise data centers needing a tool for online data migration between heterogeneous storage environments. This add-on design allows for heterogeneous storage and data movement, but does not simplify management of the configuration. Most of these designs are implemented as external appliances sitting within the storage area network. These devices generally add more complexity to an already complex environment – introducing additional points of management within the SAN, masking value-added features of the attached arrays, and limiting performance and scalability of the consolidated SAN storage pool to a single gateway hosting the backend storage. In the iSCSI realm, simpler design alternatives make the promises of storage virtualization easier to achieve.

The most advanced storage virtualization technologies offer the ability to virtualize storage at the storage device level. This approach achieves scalable performance and capacity and reduces overall management by aggregating multiple controllers into a cooperating set of resources, i.e., virtualizing volumes not only across disks within a controller but also across storage systems in a SAN. Dell EqualLogic PS Series storage arrays are an excellent example of such a "scale-out" architectural design. Ironically, despite their elegant simplicity, such designs are virtually non-existent in the Fibre Channel

SAN world, primarily due to architectural constraints inherent in Fibre Channel network deployments.

Devices in a Fibre Channel SAN are attached to the network via a World Wide Name (WWN), a physical port address specifically assigned and encoded within the device hardware/firmware. World Wide Names are assigned for all devices within the SAN, including each physical port within each host server and each physical port within the storage controllers. Furthermore, data paths between the hosts and the storage array are statically set when the host is added to the SAN.

Herein lies the inflexibility of Fibre Channel SANs. The topology of the SAN is intrinsically hard coded into the environment, making changes within the infrastructure burdensome, time consuming and error-prone. Expansion of resources in the host, fabric or storage layer propagates changes throughout the infrastructure, causing intrusive downtime to the applications and infrastructure.

By contrast, TCP/IP networks support virtual addressing and dynamic routing, whereby paths through the network are not statically defined. With DHCP, addresses can be dynamically assigned, or through address proxying, physical addresses can be virtualized, making it possible for devices in the network to transparently act on the behalf of other devices on the network. If the IP address of the resource being accessed is known from within the network, the entity (such as a SAN initiator) requesting access can find the resource dynamically without prior knowledge of the paths or the resource's physical address.

IP address proxying makes storage access fully virtualized across multiple EqualLogic PS Series storage controllers. In an EqualLogic PS Series group, for example, each member array has three active physical gigabit Ethernet ports. A 4-member group has a total of 12 active Ethernet ports, providing 12 Gb/second peak bandwidth. Each of these Ethernet ports is assigned an IP address. The members within the group are aware of and coordinate use of these multiple IP addresses and the underlying physical resources.

External access to the group by all hosts within the SAN is exclusively addressed via the group IP address, a unique IP address that transcends all the underlying member IP addresses. The only IP address known by the attached hosts in the SAN infrastructure is the group IP address. All I/O's may be serviced by any physical gigabit Ethernet port.

Shielding host systems within the SAN from being aware of the physical topology of the storage serving its data is vitally important, particularly in a virtualized server environment. By virtualizing all physical I/O ports, a flexible storage utility is created. Pooling storage assets within the SAN is now possible, enabling a rich set of features to maximize storage utilization and effectiveness, extending beyond the limitations of a single, physical end-to-end connection to a single storage controller.

Now, the host accesses the data without intimate knowledge of where in the storage utility the data is stored, providing the opportunity for data volumes to span controller boundaries and exploit the combined resources of multiple controllers. Controllers coordinate amongst themselves to balance workloads and optimize storage resource utilization.

Changes in storage infrastructure are now seamless and have no impact on host connectivity or data access. Controllers can scale the iSCSI-based infrastructure non-disruptively, automatically applying the resources of the additional controllers and disks to extend performance and capacity. They can seamlessly integrate multiple tiers of storage in the utility, providing for automatic placement of data within a volume to the appropriate tier, based on the volume's access patterns.

A flexible storage architecture as described above is particularly appealing in a virtualized server environment as the "scale out" architectures of both server and storage infrastructures uniquely complement each other, providing an end-to-end virtualized infrastructure. Deployed together, IT managers achieve an infrastructure that is particularly adaptable to changing business requirements by providing non-disruptive workload migration and balancing across all physical resources--both storage and server — as well as online resource expansion that immediately apply to pre-existing server and storage workloads, without intervention and tuning.

Dell EqualLogic Virtualized SANs

A virtualized iSCSI SAN complements — and is characteristically equivalent to — a VMware ESX Server farm.

The foundation of an EqualLogic virtualized iSCSI SAN (see Figure 2) is the PS Series storage array. Each array is fully-redundant, containing disks, multiple high-performance network interfaces, redundant controllers with mirrored battery-backed caches, and other advanced features. The disks are automatically protected with RAID (RAID 10, RAID 50, or RAID 5) and hot spares. Multiple models of arrays

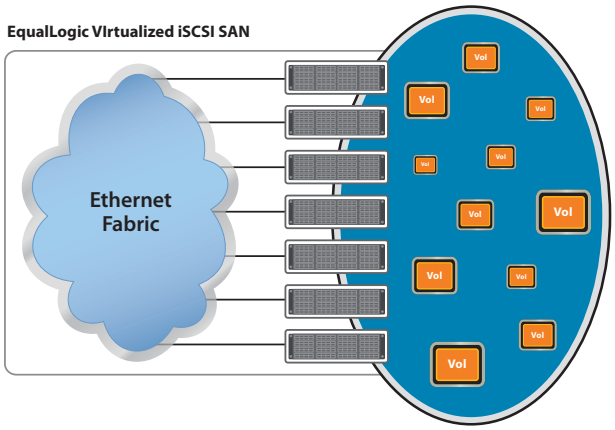

exist, leveraging high-performance 10K and 15K RPM SAS drives or more cost-effective, higher density SATA-II drives.

A PS Series storage group is comprised of a single PS Series array or multiple arrays. A group is a virtualized resource, appearing to ESX Servers as a single entity that offers network storage access to a single large pool of storage – a storage area network (SAN) composed of a single, virtually scalable, high performance storage system. Unlike a conventional SAN, in which multiple controllers appear as independently managed islands of storage, each group member in an EqualLogic PS Series SAN "cooperates" with other members to automate resource provisioning and performance optimization.

ESX Server administrators create data volumes from within the EqualLogic group storage pool. The group exports its volumes as iSCSI targets protected with security, including authentication and authorization. Upon connection, ESX servers work with EqualLogic's volumes as VMFS Datastores, which contain multiple virtual machines and virtual disks. A variation on this is a Raw Device Mapping (RDM) volume. RDM provides a mechanism for a virtual machine to have direct access to a LUN on the physical storage subsystem. A third alternative is to employ the native iSCSI software initiator of the individual guest operating system to connect directly to the SAN-resident volume.

Automatic Load Balancing

Volumes are distributed among the group's member arrays, with data placement and access continually adjusted for optimal performance as resources are added or workloads change. When an array is added as a group member, its disk space is added to the group's storage pool. Volumes are automatically re-striped and distributed across all the members of the storage pool. Controller resources are also dynamically adjusted based on the workloads being generated by the ESX Server cluster. Data and network I/O to the group are automatically load balanced across the group members' resources.

As capacity and performance requirements increase, a group can be scaled linearly in both capacity and performance – all while online. New members "learn" configuration and performance information from the group – with no manual intervention. Data and client connection load balancing occur automatically as the group scales. I/O activity is monitored, and data and network connections are adjusted as needed.

Non-Disruptive Scalability

The scalability model allows for automated, online expansion in all storage dimensions, and the PS Series architecture nearly eliminates downtime caused by expanding or managing a storage system. Because IT manager can add so easily, they only need to buy the storage necessary for today's applications, which cuts costs.

Additional benefits of the virtualized SAN environment include the transparent application of storage tiers. EqualLogic storage groups can use multiple storage tiers contained within the storage pool and automatically optimize data placement based on workload. Alternatively, storage tiers may be segmented into independent resource pools to guarantee specific resources to specific workloads, concurrently maintaining the flexibility to migrate volumes online from one tier to another, fully transparent to the host servers.

Quick, Intelligent Provisioning

The EqualLogic PS Series group makes provisioning storage fast and easy, with just the click of a button. It automates key functions for configuring, managing and scaling storage, cutting administration tasks for volume and capacity growth. Given the intelligence built into the PS Series software, decisions with respect to RAID type and data layout are made automatically during provisioning, and optimized as workload patterns for the newly provisioned storage evolve. As new resources become available, EqualLogic applies resources automatically where and when needed. By eliminating complex tasks and permitting fast and flexible storage provisioning, PS Series solutions dramatically reduce acquisition and ongoing operational costs and make enterprise class shared block storage practical for the mid-range storage market.

Thin provisioning extends EqualLogic's existing provision¬ing features. This makes the buy-as-you-grow storage model of the modular PS Series more seamless for servers and applications. Thin provisioning is an important

advanced feature that enables the automatic addition of physical capacity on demand up to preset limits.

With advanced thin provisioning, buy-as-you-grow storage management and virtualization are made seamless for servers and applications. When a volume is created, it can be sized for the long term needs of the application without initially allocating the full amount of physical storage. Instead, as the application needs more storage, capacity is allocated to the volume from a free pool.

EqualLogic's thin provisioning capability provides exten¬sive flexibility and user safety controls. These include the ability to turn thin provisioning on and off as needed for any volume. This allows users to affordably test the most suitable applications and volumes for thin provisioning, with the knowledge that they can return to "normal" provisioning online. EqualLogic's implementation of thin provisioning also provides enhanced alerts and controls – with proactive, userdefined threshold alarms and controls, administrators can depend on automatic space allocation without worrying about reaching allocation limits or unexpected depletion of physical storage.

Automated Management

The PS Series architecture is designed to simplify storage management in several ways. RAID configuration and hot sparing is automated, and dynamic storage and network I/O load balancing occurs automatically as resources and performance metrics change. No longer must administrators manually map application data to specific physical devices and controllers.

A primary benefit of this automated, virtual storage system is that storage administration remains stable as the group is expanded. EqualLogic PS Series group manager, an intuitive, single-pane administrative console, performs all administration at the group level. Because the group is managed as a single logical system, the ongoing operational costs of storage management remain fixed, even as storage grows.

A Flexible, Scalable, Virtualized Environment for Servers and Storage

iSCSI is a key technology that delivers scalable, cost-effective, high-performance virtualized SAN environments, a perfect complement to a virtualized server environment. The storage virtualization achieved in the advanced iSCSI-based design of the Dell EqualLogic PS Series enables greater performance, scalability, ease of use, and flexibility in contrast to classic storage array architectures. As a result, both server and storage assets become fully virtualized (see Figure 3), abstracted from the physical hardware upon which they reside. Virtualized

Virtualized Servers Meet Virtualized Storage

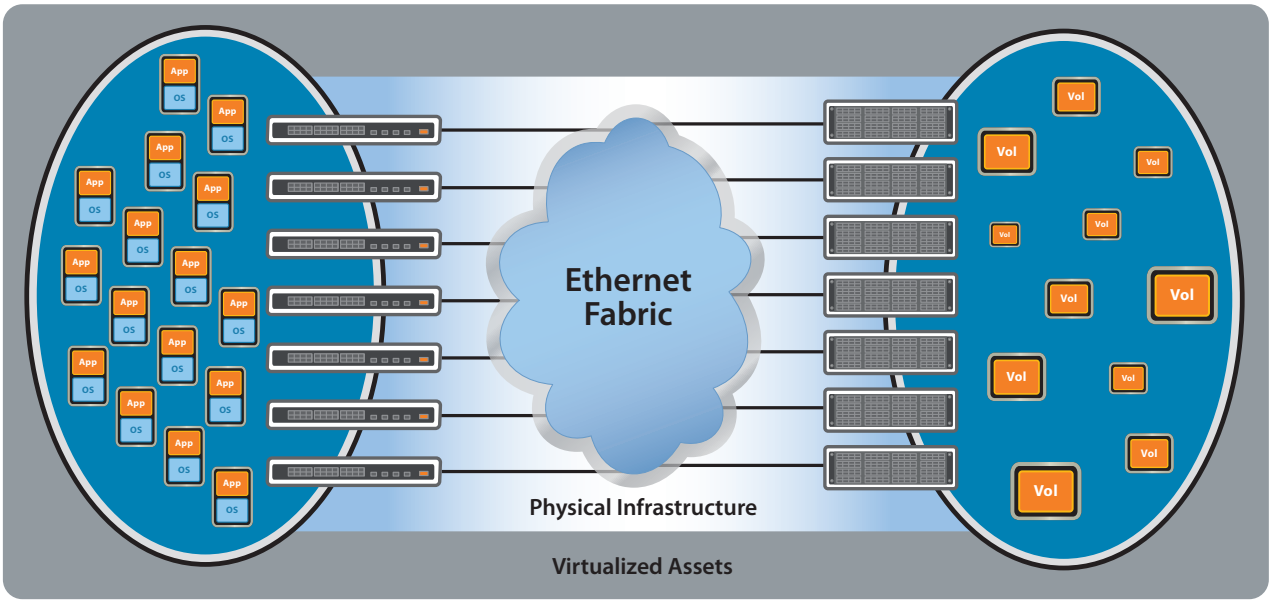

server assets leverage a pooled set of physical server resources. Similarly, virtualized storage volumes leverage a pooled set of physical storage resources.

By combining virtualized server and virtualized storage technology, the Dell EqualLogic PS Series provides a simple, flexible IT infrastructure is, resulting in a comprehensive solution with a common set of benefits — increased IT flexibility, lower total cost of ownership and reduced complexity. Key features include:

- Aggregation of virtualized assets on consolidated hardware – Operational procedures and best practices can be standardized and consistently applied to both storage and server assets (both physical and virtual) resulting in a more resilient infrastructure with greater resource utilization, higher levels of service, and enhanced protection of information assets.
- Simple, centralized management – Centralizing management through intuitive, graphical management tools accessible from anywhere on the network, provides greater management efficiencies, because there is a comprehensive view to provision, monitor, and manage the entire virtual¬ized infrastructure.
- Flexible, fast deployment of virtualized resources – Organizations can quickly adapt to changing and growing business needs by reducing the time to provision and deploy new applications via quick provisioning methods, available for both servers and storage.
- Online, non-disruptive resource re-allocation and expansion – As workflows and business priorities change, the solution can re-allocate both storage and server resources online, with no disruption to operations. In addition, it can expand physical resources expanded easily, without downtime. The solution automatically rebalances workloads d across these newly available resources without disruption to applications.
- Common IP network-based infrastructure – Basing all operations on IP networking, including the interconnect for client access, inter-server communication, storage access, and off-site data replication, simplifies the IT environment. Leveraging the organization's inherent IP networking expertise results in lower training and ongoing management costs.
- Enterprise-class resiliency – Building redundancy into the physical server, network, and storage architecture, as well as component failure detection and failover software implemented within each layer of the infrastructure, enhances overall reliability, availability and service levels.
- Advanced data management and disaster recovery – A rich set of server and SAN-based data protection tools ensure the organization's critical assets are protected and immediately recoverable at the local or a remote site.

Summary

Virtualized iSCSI SANs are changing customer experiences of how simple an enterprise storage infrastructure can be to deploy, manage and grow. They are uniquely positioned to inspire broad adoption of virtualized server technologies by reducing technical complexity and cost barriers imposed by classic storage area network architectures, without compromising the performance, scalability and resiliency requirements of a virtualized IT infrastructure. As a result, organizations of all sizes are now deploying EqualLogic PS Series virtualized iSCSI SANs as the backbone of their virtualized infrastructure.

VMware Infrastructure makes more efficient use of a company's existing resources, reducing costs and helping companies respond to business needs faster with a virtualized IT infrastructure. Characteristically, both VMware Infrastructure and storage virtualization provide key features for reducing complexity and increasing flexibility, including virtualized asset management, incremental online growth, workload migration and load balancing. The complementary affects of Dell EqualLogic virtualized iSCSI SANs and VMware virtualization technologies make possible a simple, cost-effective, and dynamic enterprise-class IT environment.

CUSTOMER CASE STUDY: COSTAR GROUP, INC.

Costar Group, Inc., based in Bethesda, MD, is recognized as the leading provider of information services to commercial real estate professionals in the United States and the United Kingdom. On a daily basis, real estate professionals throughout the United States and the United Kingdom rely on Internet access to Costar's extensive database that includes information on more than 2.2 million commercial properties, 780,000 properties actively being marketed for sale or lease, and more than $460B of properties for sale.

With annual revenues generated exclusively from online subscription and services revenue approaching a run rate of $200M, the company's success and future growth are critically dependent on a robust, scalable IT infrastructure. With more than 1,300 employees accessing and updating the online services portfolio and 75,000 subscribers using their online services, the IT organization is always looking for ways to enhance and simplify their burgeoning infrastructure. Costar turned to virtualization to reign in an expanding IT footprint and reduce spiraling management costs.

Costar has built a virtualized IT environment based on VMware Infrastructure 3 and an EqualLogic PS Series virtualized iSCSI SAN as a key element of their infrastructure.

By leveraging both server and storage virtualization, Costar has architected a solution that meets their goals to reduce IT footprint, increase responsiveness for provisioning new IT assets, simplify the ongoing management of both server and storage assets, and seamlessly scale their infrastructure over time.

VIRTUALIZED SERVERS

The production environment includes 8 ESX Servers based on dual & quad-processor, dual-core Opteron-based servers each with 24 to 32 GB's of RAM in support of a phased rollout of more than 200 virtual machines. Three additional ESX Servers are dedicated to test and development. QLogic QLA4052C iSCSI HBAs are configured in each of the ESX Servers to enable SAN boot capability, Jumbo Frames support, and reduced I/O processing overhead.

At present, 115 virtual machines are in production, hosting a variety of Windows Server 2003R2-based applications including SQL databases, time accounting, source control, helpdesk applications, data creation and management tools, and internal back office processing. VMware's VMotion and High Availability (HA) features are actively used for workload balancing, proactive maintenance, and increased server availability.

VIRTUALIZED STORAGE

The storage backbone of the VMware infrastructure is an EqualLogic virtualized iSCSI SAN, connected with the ESX Server environment over a fully redundant gigabit Ethernet fabric. Comprised of 5 EqualLogic PS Series arrays, the EqualLogic group contains 80 high performance, 10K RPM Serial-Attached SCSI (SAS) disks configured as RAID-50 with hot spares, for a total usable capacity more than 15 terabytes.

Currently, the 11 ESX Servers access nearly 4 TB's of provisioned storage, configured as 8 VMFS file systems and 11 ESX boot volumes, all resident on the EqualLogic SAN. Each VMFS file system's data volume is auto-load balanced across the 5 member arrays and all the disk drives in the SAN. The aggregation of workloads from across 100 virtual machines in the 8-member ESX Server cluster produces a random workload profile that particularly benefits from the multi-controller architecture and aggregate performance of the large number of fast spinning drives.

FUTURE DIRECTIONS

Future expansion of Costar's virtualized infrastructure will include the rollout of an additional 100 virtual machines, including their Exchange environment and additional SQL Server databases. Costar is also planning to add a second tier of virtualized storage; the EqualLogic group will be expanded online to include two additional SATA-II-based PS Series arrays as a secondary storage pool. The EqualLogic instant cloning and online volume migration features will be used to quickly generate test and development environments from production data.

Additional focus will also be placed on disaster recovery processes; a second EqualLogic group, comprised of 2 PS Series arrays, is already in place as the disaster recovery target for the VMFS file systems. Using EqualLogic's auto-replication feature, the production VMFS volumes are being replicated to the EqualLogic group at the corporate DR site.

VMware, Inc. 3401 Hillview Ave Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

For more information regarding Dell EqualLogic solutions, contact us at 888-579-9762 ext 7792 or visit www.equallogic.com.
For more information on EqualLogic solutions in a VMware environment, visit www.equallogic.com/vmware.
For a live, web-based product demonstration, visit www.equallogic.com/demo.
110 Spit Brook Road, Building ZKO2, Nashua, NH
03062 Tel 603.579.9762 / Fax 603.579.6910
www.equallogic.com