

STORAGE OPTIONS FOR MICROSOFT[®] EXCHANGE[™] ENVIRONMENTS

MANAGEMENT SERIES

Electronic communication services are considered a critical business application that dramatically impacts the daily operations of any organization. As e-mail applications are heavily dependent on storage resources, it is useful to explore the tight integration between Microsoft Exchange environments and storage platforms.

This paper provides an overview of various storage options available for Microsoft Exchange environments, and describes the challenges, advantages, and disadvantages of each.

When developing Windows Server 2003, Microsoft made a conscious effort to provide enhanced support for networked storage solutions, including new services and drivers designed specifically for higher performance and fabric management of SANs.

In the past decade, e-mail and electronic messaging have become the backbone of communication and collaboration within and among organizations, with Microsoft Exchange by far the most often deployed messaging environment in the corporate marketplace. This mass adoption has resulted in more rapid and efficient communications that organizations have quickly embraced, enhancing business processes and productivity. Although revolutionary in its impact, this mass adoption of e-mail communication has affected business processes and created new IT challenges, primarily focused on governing and keeping pace with the rapidly growing dependency on e-mail.

As a primary vehicle for sharing information and transmitting documents, e-mail has become drastically more storage resource intensive. Additionally, the dependence on e-mail communication in business has made it essential to have highly available and optimally performing Microsoft Exchange environments that can not only maintain and ensure business continuity, but also provide fast data recovery in case of downtime or disaster. Further more, the state of today's compliance environment often requires organizations to keep electronic communications readily available for extended periods of time – this drives the e-mail capacity requirements higher than ever. As a result of these fundamental shifts in e-mail solution requirements, the supporting IT infrastructure – namely, the storage – must be high performance, flexible, scalable, and above all, reliable, to meet today's demands as well as those of tomorrow.

Microsoft Exchange, the dominant corporate e-mail application, continues to proliferate based on its market leadership, simplicity, and close integration with business applications from Microsoft and other vendors. Microsoft is working towards integrating Exchange with other collaboration and voice applications to further improve business processes, which will help Exchange to maintain its momentum in the corporate messaging market. These new integrations are driving organizations to rethink their storage strategies. A virtualized, scalable, easy-to-manage storage infrastructure is required to support both increasing application integration and growing storage capacity demands.

THE ROLE OF THE WINDOWS OPERATING SYSTEM

Before selecting a storage architecture for your Microsoft Exchange infrastructure, it is important to consider the role of the operating system (Microsoft Windows) on which your Exchange server will reside, and how well your storage solution integrates with key Windows features.

Microsoft Windows Server 2003 R2 can help organizations reduce storage management costs while improving storage efficiency. With continuing improvements around data availability and storage connectivity as well as the recent introduction of new storage management consoles, Windows Server continues to grow as a more integrated server and storage infrastructure.

Enhanced, built-in support for Storage Area Network (SAN) hardware from multiple vendors allows easier, more reliable access to organizational data, without the need for third-party tools. Microsoft Windows Server 2003 R2 includes several new features to help organizations better manage corporate data, improve storage utilization.

When developing Windows Server 2003, Microsoft made a conscious effort to provide enhanced support for networked storage solutions, including new services and drivers designed specifically for higher performance and fabric management of SANs. While some of these SAN management capabilities have been back-ported to support Windows 2000 Server, the majority are unique to Windows Server 2003. These services include:

- Volume Shadow Copy Service (VSS), which enables current state “snapshot” style backups of volume, folders, or files
- Virtual Disk Service (VDS), which enables software applications to configure and manage SAN arrays independent of specific implementations via a hardware provider model
- Enhanced disk management capabilities such as the ability to grow volumes (using VDS)
- Support for remote SAN boot, flexible volume mounting, and an enhanced driver model for SAN deployment scenarios
- Multi-path I/O (MPIO) support for high availability and load balancing

The limitations of DAS, especially the lack of scalability and availability, expose Microsoft Exchange services to unacceptable levels of planned and unplanned downtime and make it very expensive to scale and keep up with the growth in demand.

While it is not mandatory to integrate all of the features and functionality offered by Microsoft Windows into your storage solution, it is important to choose a storage vendor that is dedicated to integrating the ongoing improvements led by Microsoft at the server and application levels.

STORAGE OPTIONS

Traditionally, Exchange and mail servers have been deployed using Direct Attached Storage (DAS), which initially satisfied the requirements of those servers. Although the performance was acceptable to properly run e-mail applications, the downtime problems due to drive failures and regular maintenance drove system administrators to look for different technologies to overcome these restrictions. Administrators turned to new storage implementations to take advantage of the scalability, availability, or performance improvements of technologies such as Network Attached Storage (NAS) and Fibre-Channel (FC) SANs – which, despite higher prices and complexity, offered a better alternative for administrators. Today, SCSI over Internet Protocol (iSCSI) SAN technologies offer more compelling advantages for e-mail services, delivering the simple-to-implement, cost-effective, scalable solution for which administrators had been searching.

DIRECT ATTACHED STORAGE


DAS is defined as a small collection of disk drives (typically 2-12) attached to a single server via a dedicated connection. DAS attributes its roots to the era of computing before the networking revolution. In the 1990s, with the advent of cost-effective, application-based servers, companies deployed islands of these servers, each dedicated to a different business application. At the time, infrastructure design was not concerned with the future needs of data sharing or information management; the focus was on applications and servers, not data and storage. Over time, these deployments revealed the limitations of DAS, including:

- **Single Server Access:** Storage is accessible by the hosting server only.
- **Many 'Single Points of Failure':** Any component failure in the server makes the data inaccessible, resulting in unpredictable service levels and unscheduled business disruption.

- **Labor-Intensive Management:** Each DAS unit presents an isolated system to be individually managed, configured, and backed up.
- **Finite Scalability, Costly Upgrades:** Hardware limits present a finite number of disks, arrays, and cards that can be added to a single system. Also, adding DAS capacity often requires hours of system downtime and costly staff time.
- **Inefficient Utilization:** These problems force administrators to provision excess storage in advance, with the hope of avoiding costly upgrades, resulting in poor asset utilization and greater than necessary initial investments.

These limitations, especially the lack of scalability and availability, exposed Microsoft Exchange services to unacceptable levels of planned and unplanned downtime and made it very expensive to scale and keep up with the growth in demand.

Peer Storage Utilization in DAS Environments



CONSOLIDATION WITH NAS FILERS

Often referred to as “filers”, NAS systems are defined as servers or appliances that export file systems over the network. Workstations and servers use a network file protocol (most commonly, CIFS or NFS) to access shared files stored on disks within the filers. These filers have several advantages – they are relatively easy to setup and use, and they leverage the standard Ethernet infrastructure to connect clients to the storage. Many NAS configurations are available, for environments ranging from enterprise datacenters to distributed workgroups.

Despite its benefits, NAS has several drawbacks and leaves some basic problems unsolved. Network file protocols can slow performance compared to more efficient

block protocols. Even more important, storage designed to provide file access (as NAS is) is not designed for running databases. In addition to performance hits, traditional NAS systems suffer from the same scalability problems associated with DAS. The filer protocols are often incompatible with certain common applications – in fact, Microsoft does not support Exchange deployments that use NAS as the backend for mailboxes.

A BETTER APPROACH FOR EXCHANGE: STORAGE AREA NETWORKS

SANs offer a block-level protocol (SCSI) for clients to access data over the network. This block-level protocol is then encapsulated by a network protocol such as Fibre-Channel or iSCSI. Unlike the network file protocols that NAS uses, these block-level protocols provide complete application compatibility with Exchange, and deliver performance that typically exceeds DAS.

FIBRE CHANNEL STORAGE AREA NETWORKS

Fibre-Channel is the most widely deployed SAN technology. FC was originally developed to solve the cabling problems associated with connectivity and distance. Eventually, the industry expanded the technology for SAN application, and FC today offers an extremely robust shared-storage solution. However, despite its benefits, the number of actual FC deployments for Exchange remains relatively modest, due to several limitations of FC, and once again several key problems left unresolved.

First, it is very expensive to deploy and grow FC SANs. In a Fibre-Channel environment, the adaptors, switches, and cabling used to link SAN components together are very costly, and well out of reach for most of the mainstream market. In addition, FC SAN deployment and maintenance is complicated, and requires special IT expertise that DAS and NAS do not. Because of the added complexity of FC SANs, expensive software is often required to automate and simplify operations, particularly for large deployments. Severe interoperability issues between equipment vendors compound these problems.

Second, although FC SAN addressed some of the scalability issues shared by NAS and DAS, it brought new

issues related to complexity and costly capacity upgrades. When additional storage is needed, the administrator must either add storage devices (and accompanying complexity) to the network, or upgrade to a larger, more expensive device whose installation requires system downtime. To avoid these problems, IT departments must either buy and install expensive management and virtualization software, or continually over-buy storage capacity.

Other than for very large Exchange deployments, most organizations have failed to see the cost justification for running Exchange on a FC SAN. Although FC SAN presents a good platform for Microsoft Exchange environments, its price tag makes it difficult for many organizations to switch to an FC SAN storage environment.

THE BEST OF BOTH WORLDS: iSCSI

iSCSI combines the venerable SCSI block-level protocol with Ethernet, the most widely used networking technology. The result is a highly scalable SAN technology that leverages Ethernet's affordable and familiar infrastructure. With the adoption of the standard by the Internet Engineering Task Force (IETF) and strong vendor support from industry leaders such as Microsoft and Cisco, iSCSI has rapidly matured into a practical and popular SAN technology – and specifically the right storage technology for Microsoft Exchange.

How iSCSI Works

iSCSI serves the same function as the Fibre-Channel network protocol, transporting SCSI commands and responses back and forth between the servers hosting Exchange and the storage. Whether FC or iSCSI is used, the result is the same: the operating system sees SCSI disk devices.

With FC, the hardware layer and the fabric network are unique to storage, so special components, Host Bus Adapters (HBAs) and switches are always required. With iSCSI, an existing Ethernet network can be utilized, thus requiring no additional networking infrastructure. iSCSI can be implemented using HBAs that attach to standard IP switches, or with the Windows iSCSI software initiator that presents itself to the server like an HBA, but instead

uses existing NICs to attach to the switches. The vast majority of servers, and nearly all servers running Exchange, use standard NICs with iSCSI initiator software. Even a large, 8,000-user Exchange environment requires only about half the speed of a single GbE NIC port. In addition, most servers have multiple ports and can leverage Microsoft's MPIO for higher availability and performance if needed.

The Economics of iSCSI for Exchange

The opportunity to use a standard Gigabit Ethernet NIC to connect Windows 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 there is no additional cost and administrators are not burdened with installing and configuring complicated adapter cards.

Together, standard NICs and the lower cost of Ethernet switches offer a compelling cost savings over Fibre-Channel while maintaining 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.

Table A: Comparative Table of Storage Options

	Low TCO	Ease of Use Setup	Ease of Management	Redundancy	Embeds into Ethernet Network	High performance	Comprehensive Protection	Integration with Windows
DAS	X	X	X		X			
SAN (Fibre Channel)				X		X	X	X
SAN (iSCSI)	X	X	X	X	X	X	X	X

The Administrator Dilemma

While iSCSI offers great advances for the network portion of the SAN, administrators still face storage challenges. First, Ethernet IP expertise does not translate into the skills required to successfully manage storage systems. The traditional storage administrator requires a rarified level of knowledge, from RAID geometries to application I/O workloads. Indeed, this expertise will be taxed on a routine basis for the regular deployment and maintenance of storage systems.

Table A provides an overview of how these storage technologies compare. While Exchange has traditionally been deployed on each of these storage technologies, the benefits of running Exchange on an iSCSI SAN are clear.

PS Series Solution: Self-managing, Self-scaling, Self-healing iSCSI

PS Series technology was developed by EqualLogic® to take storage management to the next level by driving intelligence down to the storage array. All of the tasks associated with daily storage management, load balancing, and provisioning are handled by PS Series arrays transparently to the administrator -and all are accomplished while the system is online, through a common intelligent interface. In addition, the PS Series architecture enables almost limitless scaling, to well over 100TB of data.

PS Series disk arrays automate tasks internally, with minimal intervention. Multiple arrays can work together, presenting a single storage view to both administrators and applications consuming the storage. Scalability issues are eliminated – the task of balancing workloads across disks and arrays is delegated to the systems themselves rather than to the administrator, and adding arrays to an existing deployment requires minimal effort.

Highest Performance Mid-range Storage Solution for Exchange

The Microsoft Exchange Solution Reviewed Program (ESRP) was developed by Microsoft to provide a common testing framework for storage vendors to offer information about their products' support of Microsoft Exchange 2003 Server software. EqualLogic has published two reports on

The ESRP tests results position the PS Series as the fastest mid-range storage array to serve Microsoft Exchange Server deployments.

Microsoft's ESRP Web site showing an EqualLogic SAN consisting of PS3800XV storage arrays supporting 10,000 and 20,000 Microsoft Exchange users. These tests results position the PS Series as the fastest mid-range storage array to serve Microsoft Exchange Server deployments.

Effective Growth Management

PS Series arrays address the Exchange infrastructure growth problem by automating all aspects of storage expansion. Arrays can be added to existing deployments without any disruption to Exchange users, and the new storage resources are immediately available to the virtualized storage pool. These resources include more than just additional capacity – they include more disk spindles, increased controller cache, and more GB Ethernet ports for additional network bandwidth. Each of these enhances overall I/O performance to support more Exchange users.

How PS Series Solutions Meet Escalating Service Expectations

Numerous virtualization strategies on the market today offer some limited forms of growth management. However, none do so while providing the enterprise services to match the increasing importance of online information. Microsoft Exchange Server is a classic example of changing priorities, as preserving e-mail is now crucial to daily business operations. IT departments must meet these new challenges head on. High availability is first on the list of desired service features – information must be available around the clock, without vulnerability to equipment failure or system downtime. PS Series arrays use RAID for complete reliability, but that disk protection is just one piece of the reliability/redundancy equation. Redundant components eliminate any single points of failure. These are critical requirements for any serious storage deployment.

All EqualLogic products include, at no additional charge, all of the enterprise class software features that are required for data center deployments. PS Series solutions offer Snapshot capability, an effective backup and recovery solution that minimizes risks. Snapshots create an instantaneous copy of a volume at a particular point in time. Subsequent changes to the volume are tracked, and

only the changes are stored. The result is that at any time, the volume can be "rolled back" to a particular snapshot; the base volume is restored to the exact state of the volume when the snapshot was taken. PS Series arrays can also remotely mirror volumes at significant physical distances through volume replication across IP networks. The PS Series support of VSS provides a clean snapshot and backup process that insures data integrity and significantly reduces the backup window.

Microsoft strongly recommends that the original data be preserved in case it is needed later for data extraction or analysis. Most storage solutions that have instant rollback permanently wipe out the most recent data and any snapshots newer than the chosen restore point. This makes it unattractive to use instant restore, so users generally wait for the administrator to manually copy the data. EqualLogic instant rollback preserves all the snapshots and preserves an additional snapshot of the original data automatically.

EqualLogic adds the capability to instantly make a full, permanent volume out of an existing snapshot. This makes it easy to preserve crucial data for future analysis even if snapshot space is later consumed by changes to the volume.

PS Series snapshot protection is further distinguished by supporting large numbers of copies and multiple retention schedules, allowing snapshots to be made frequently while also keeping a history. If problems are discovered quickly, there can be many closely spaced recovery points available; if problems are discovered later, older recovery points will still be there.

The PS Series is not just designed to meet your Exchange needs, but is also the ideal SAN to consolidate all of your server and storage resources. With EqualLogic's tiering and pooling capabilities, you can maximize Exchange performance with higher performing drives and RAID sets, but also cost effectively store less performance-critical data on drives and RAID sets designed for capacity utilization. In addition, all your storage resources are managed through a single, easy-to-use interface. This also allows organizations to provide different service levels to different Exchange users or group of users.

The PS Series provides a high performance, intelligent, full-featured storage solution that simplifies the overall storage management experience.

The PS Series superior rate of Exchange user support per disk and the all-inclusive features significantly reduce the system's acquisition cost. In addition, the PS Series' automation, load balancing, and ease of use and management reduce administrative expenses throughout the lifetime of the system. Plus, the increased levels of service offered by PS Series arrays deliver both lower Total Cost of Ownership and higher Return On Investment.

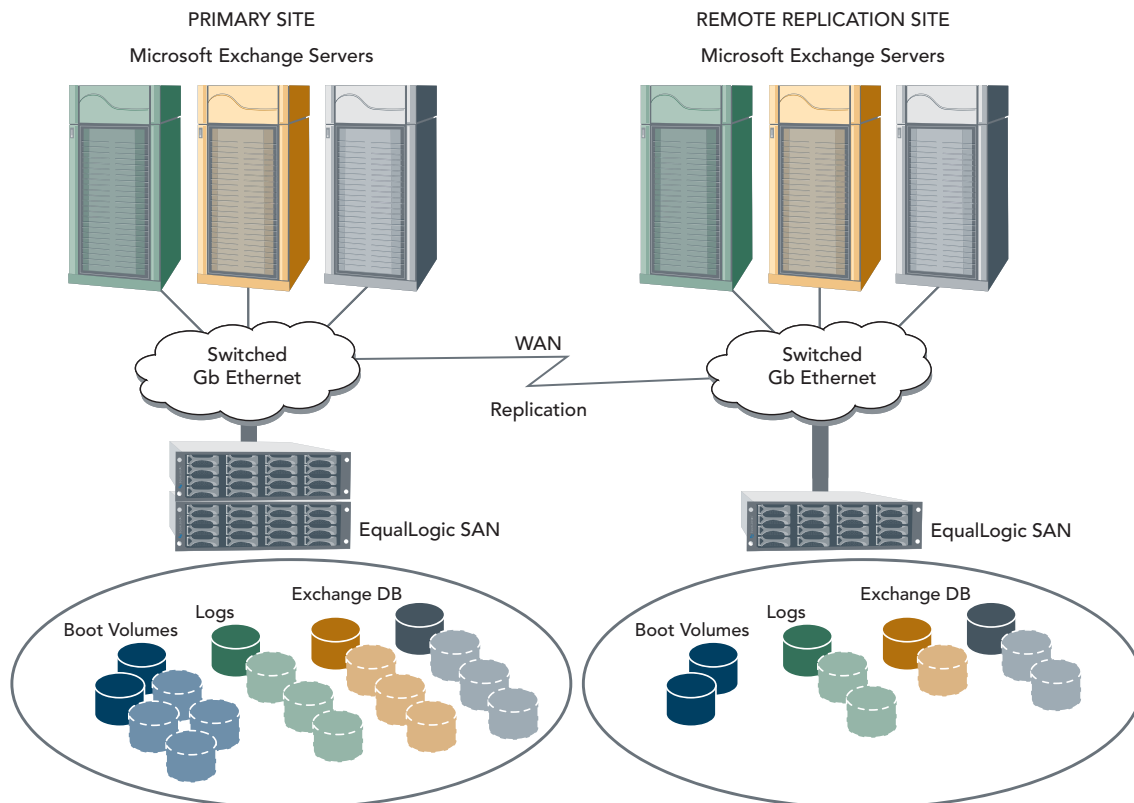
PS SERIES SOLUTION: IDEAL PLATFORM FOR EXCHANGE

EqualLogic's "next evolution in storage design" leverages core storage technologies and builds on the innovations of storage over IP as well as breakthroughs such as virtualization. The PS Series provides a high performance, intelligent, full-featured storage solution that simplifies the

overall storage management experience by better addressing the challenges that launched the storage networking revolution: managing explosive data growth and meeting escalating service requirements without adding complexity or increasing cost.

The deployment of the PS Series as the storage platform for Microsoft Exchange Server provides a comprehensive yet flexible implementation that is fast and matches the high performance expectations of e-mail services. Because of its seamless scalability, IT administrators can easily adapt to future growth requirements of more and larger mailboxes as well as greater number of users. Standard high availability and disaster features ensure business continuity and disaster protection levels that are expected in today's Exchange deployments.

PS Series Local and Remote Data Protection for Exchange Environment.





110 Spit Brook Road, Building ZKO2, Nashua, NH 03062
Tel 603.579.9762 / Fax 603.579.6910 / www.equallogic.com