

The Next Evolution in Storage Design: *Intelligent iSCSI Solutions*

Analysts estimate that business data is increasing at a 68% compound growth rate with storage pools nearly doubling every two years. Even in the face of such growth, IT organizations are expected to raise the bar on standards of availability, reliability and overall functionality without additional cost or complexity.

In order to hit this target, businesses need to implement the tools to effectively manage and maximize their storage resources.

Intelligent iSCSI solutions, the next evolution in storage design, are such a tool for companies to meet both business and technology objectives.



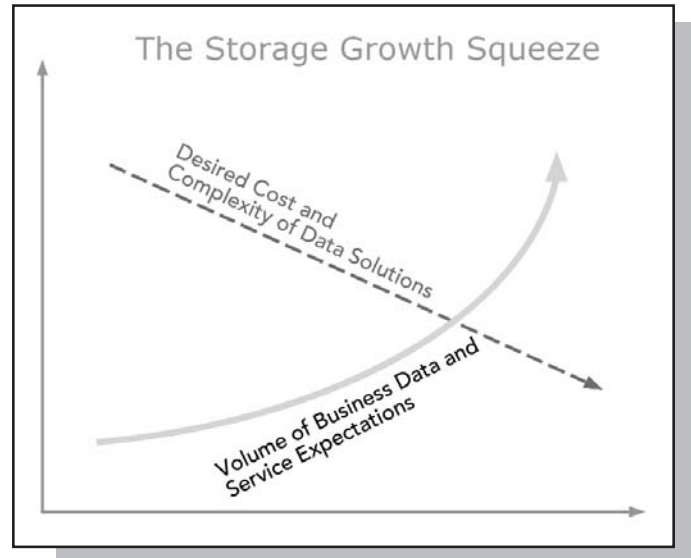
The Storage Growth Squeeze

Today, IT Professionals face two pressure-packed data storage challenges:

1. Managing the prolific growth of data, and
2. Meeting escalating service expectations

Businesses in every industry are accumulating data at an average compound growth rate of 68%, and much of this data is vital to operations. Storage pools continue to nearly double every two years. At the same time, business requirements demand that IT deliver increasing levels of availability, reliability, and functionality.

To make matters worse, IT departments are expected to manage these growing data volumes and high level services without additional costs or complexity. Storage administrators end up in the middle, squeezed by these competing requests.



These problems are not new, and many industry innovations have tried to address them. However, until recently, a complete solution remained out of reach. IT professionals still struggle to answer key questions, such as: "How can managers get more done with less?" "What options are available?" "Are there any truly new solutions?"

First, There Was DAS

"Direct-attached storage," or DAS, is defined as a collection of storage disks attached to a single server via a cable, and is the most common method of data storage. DAS attributes its roots to the era of computing before the networking revolution. In 1990s, with the advent of cost effective application-based servers, companies deployed islands of these servers, each addressing a different business problem. 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.

Single Server Access:

With the storage directly attached to one server, only that server can access the data. Data sharing is impossible.

Many 'Single Points of Failure':

Any component failure in the server makes the data inaccessible. If a cooling fan fails, users cannot access data even though the storage is functioning perfectly. These risks result 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. The complexity of managing DAS grows exponentially with the number of DAS units. A full-time administrator is required for every 1.5 to 5.0 TB of DAS storage.

Over time, deployments revealed the limitations of DAS.

Finite Scalability, Costly Upgrades:

Hardware limits present a finite number of disks, arrays, and cards that can be added to a single system. In addition, adding DAS capacity requires hours of system downtime and costly staff time.

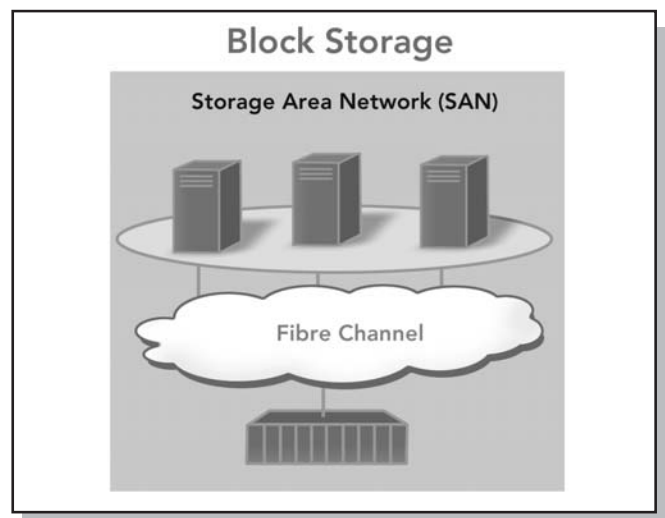
Inefficient Utilization:

These problems force administrators to provision excess storage in advance, with the hope of avoiding costly upgrades. As a result, asset utilization in DAS systems is poor, with only 20-30% of storage capacity utilization realized on average. In addition, dynamic information systems are unpredictable - administrators do their best to forecast where they expect growth, but often their predictions are incorrect. Because DAS units are completely isolated from each other, administrators cannot correct imbalances in storage use between servers. The result is a costly waste of storage resources.

The Consolidation Trend

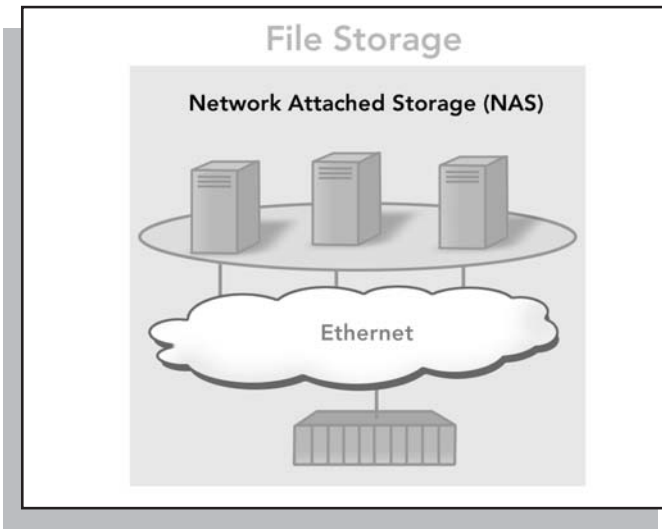
To overcome the shortcomings of DAS, the storage industry developed a new solution: storage consolidation. This infrastructure design offers a way to share storage resources using one of two technologies: Network-Attached Storage (NAS) and Storage Area Networks (SANs).

Simply put, these shared storage technologies make a collection of storage resources available to multiple users. Storage is migrated from many isolated DAS systems onto a



single, large shared-storage system; the servers then access the storage over the network.

Consolidating data onto a shared-storage system solves some of the problems associated with DAS. However, NAS and SAN environments leave many problems unresolved, and create new challenges for IT.



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 (CIFS and NFS are the most common) to access shared files stored on disks within the filers.

These filers have several advantages - they are relatively easy to 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 still leaves some basic problems unsolved. Chiefly, the filer protocols are often incompatible with certain common applications. For example, Microsoft Exchange™ deployments that use NAS as the backend for email mailboxes are not supported. For these applications, the performance limits caused by network file protocols are unacceptable.

More problematic than performance, however, NAS systems suffer from the same scalability problems associated with DAS. NAS has a finite capability to scale, and upgrades are cumbersome and costly. Once a filer runs out of disk space, the administrator must either add another filer or replace the old filer with a larger model. Neither option is desirable. Adding new devices adds management complexity, and upgrading to a larger device is expensive and requires downtime. Over time, the performance and scalability problems of NAS compound the management problem, as administrators are required to constantly rectify imbalances and bottlenecks across multiple, disparate NAS devices.

NAS systems suffer from the same scalability problems associated with DAS.

A Better Approach: Storage Area Networks

Storage Area Networks (SANs) offer a block-level protocol (SCSI) for clients to access data over the network. This block-level protocol is embedded into a network protocol such as Fibre-Channel (and more recently, IP with iSCSI). Unlike the network file protocols that NAS uses, these block-level protocols provide complete application compatibility, plus performance results that match and often exceed those of DAS.

Fibre-Channel (FC) is the most widely deployed SAN technology. FC was originally developed to solve the cabling problem associated with parallel SCSI disk arrays. 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 remains relatively modest. This is due to several new problems created by FC, and once again several key problems left unresolved.

First, it is very expensive to deploy and grow FC SANs - prohibitively so for most organizations. In a Fibre-Channel environment, the adaptors, switches, and cabling used to link the SAN 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. Initial FC SAN deployments require complete upheaval of hardware and software, and often a revolutionary change in the datacenter's architecture. Because of the added complexity of FC SANs, expensive software is often required to automate and simplify operations, particularly for large deployments.

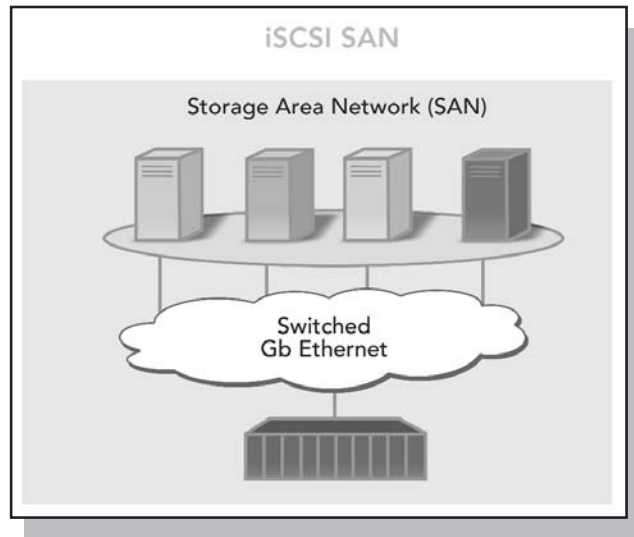
Severe interoperability issues between equipment vendors compound these problems. Conflicts between product lines offered by the same FC vendor are not uncommon. While recent improvements have been made, continuing compatibility issues mean administrators must navigate complicated matrices of supported configurations before implementing a solution.

Finally, FC SAN devices share the scalability issues of NAS and DAS - tiered scaling, plus complex 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 an expensive larger 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.

For large datacenters hosting mission critical data and applications, these problems don't preclude FC SAN implementations. A special staff of administrators can manage the huge data volumes and additional complexity, and extra staffing and equipment expenses are often justified by the economies of scale. For this reason, FC deployments have found the most success in high end, enterprise datacenters.

The Best of Both Worlds: iSCSI

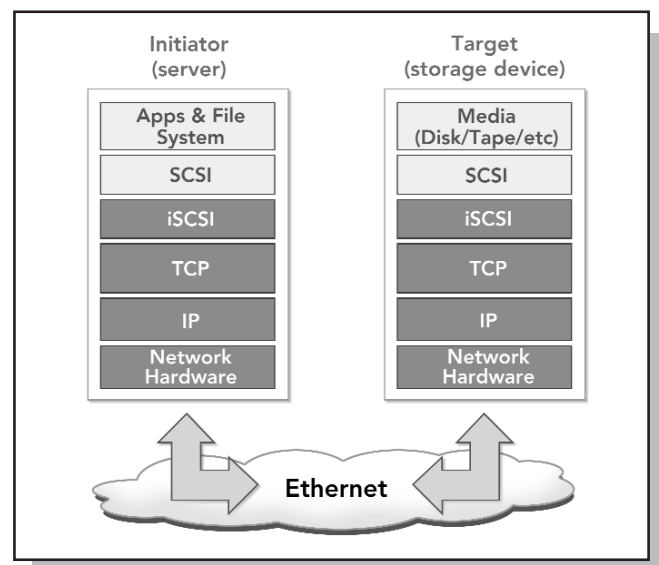
Internet SCSI, or iSCSI, combines the venerable SCSI block-level protocol with Ethernet, the most widely used networking technology. The result is a highly scalable SAN technology that leverages Ethernet's affordable and familiar infrastructure. With the adoption of the standard by the IETF (Internet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



How iSCSI Works

Utilizing 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 in packets for the TCP/IP layer. Packets are sent over the network using a point-to-point connection. Upon arrival, the protocol translates data back to SCSI. Security is provided through iSCSI authentication and virtual private networks (VPNs), as needed.

When an iSCSI initiator connects to an iSCSI target, the storage is seen by the operating system as a local SCSI device; this device can be formatted like any other local SCSI device. The process is transparent to applications, file systems, and operating systems. By consolidating storage with an iSCSI SAN, different platforms can share the same storage, greatly improving utilization and efficiency. Multi-protocol switches enable co-existence between iSCSI and Fibre-Channel SANs.



To access iSCSI storage, a server requires only an iSCSI initiator connected to a network. That initiator can be an iSCSI driver with a standard network card, or a card with a TCP offload engine (TOE) to reduce CPU utilization. Also available are HBAs that offload both TCP and iSCSI. On the target side, storage devices similarly implement the iSCSI protocol stack. The following diagram illustrates the different approaches to iSCSI initiators:

Features that storage administrators have come to expect are all available with iSCSI: booting servers from the SAN, high availability through redundant physical paths, storage over vast distances, and security. All of these features are designed into the standard, and in many cases utilize the well established solutions of IP Ethernet. In addition, interoperability is at the core of iSCSI design. Administrators can rest assured that their favorite network switches and NICs will be compatible with all iSCSI products.

The Economics of iSCSI

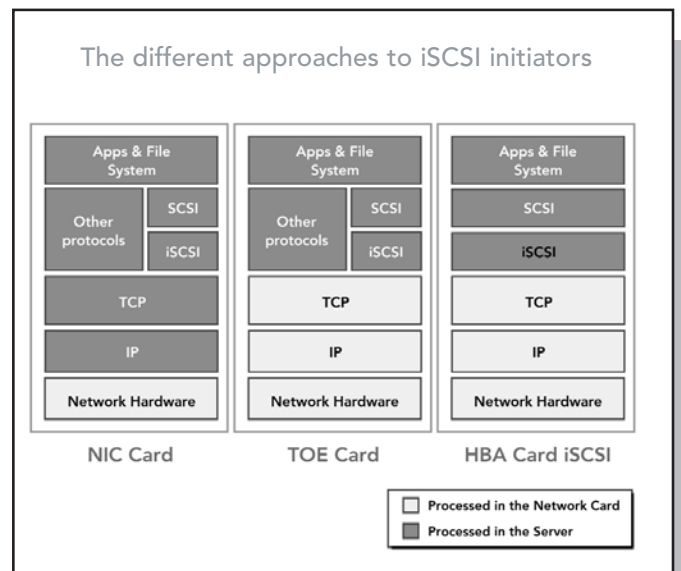
The opportunity to use a standard Gigabit Ethernet NIC to connect servers to storage makes iSCSI both simple and inexpensive. Today's high speed CPUs can run iSCSI at line speed over standard NICs using only a marginal amount of CPU. Onboard Gigabit Ethernet NICs ship standard with servers, so the cost is included and the administrator need not bother with installing complicated adaptor 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.

The Right Network, But the Same Storage Problems

While iSCSI offers great advances for the network portion of the SAN, the problem is only half solved. The storage challenges remain the same. 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.



Fundamentally, the finite scalability problem shared by DAS, NAS, and SAN storage systems remains. Disks have physical limitations in capacity and performance, and the RAID sets that span these disks inherit these fixed limits. While some storage controllers can accept expansion chassis and additional drives, eventually these systems max out. To work around these limitations, administrators overbuy, manage multiple systems, or upgrade when systems are overrun.

The "brute force" approach of overbuying is the least cost effective. On the other hand, upgrading to a larger system is not only expensive, but typically involves a serious disruption of service. Simply adding more systems to the SAN seems to be logical solution.

However, managing multiple storage systems on the SAN is a difficult proposition. Application workloads must be rerouted manually, a task that requires understanding the nuances of each disk array and the RAID sets within. Tuning and tinkering are required as workloads change over time. The problem is compounded as the storage resources themselves are shared by multiple applications.

But storage arrays cannot manage themselves, and cannot work together to solve the problems. Or can they?

Expensive software can solve the problem, but may create others: OS and hardware incompatibilities, network security issues, and vendor lock-in. Host-based software solutions often create more problems than they solve, fail to hide complexity, and fall well short of reducing costs.

Some vendors propose putting intelligence into the network itself, by placing so-called "smart" switches between the storage devices and the servers. These solutions simply add another component that needs to be managed, and they fail to remove the complexities of configuring backend storage arrays.

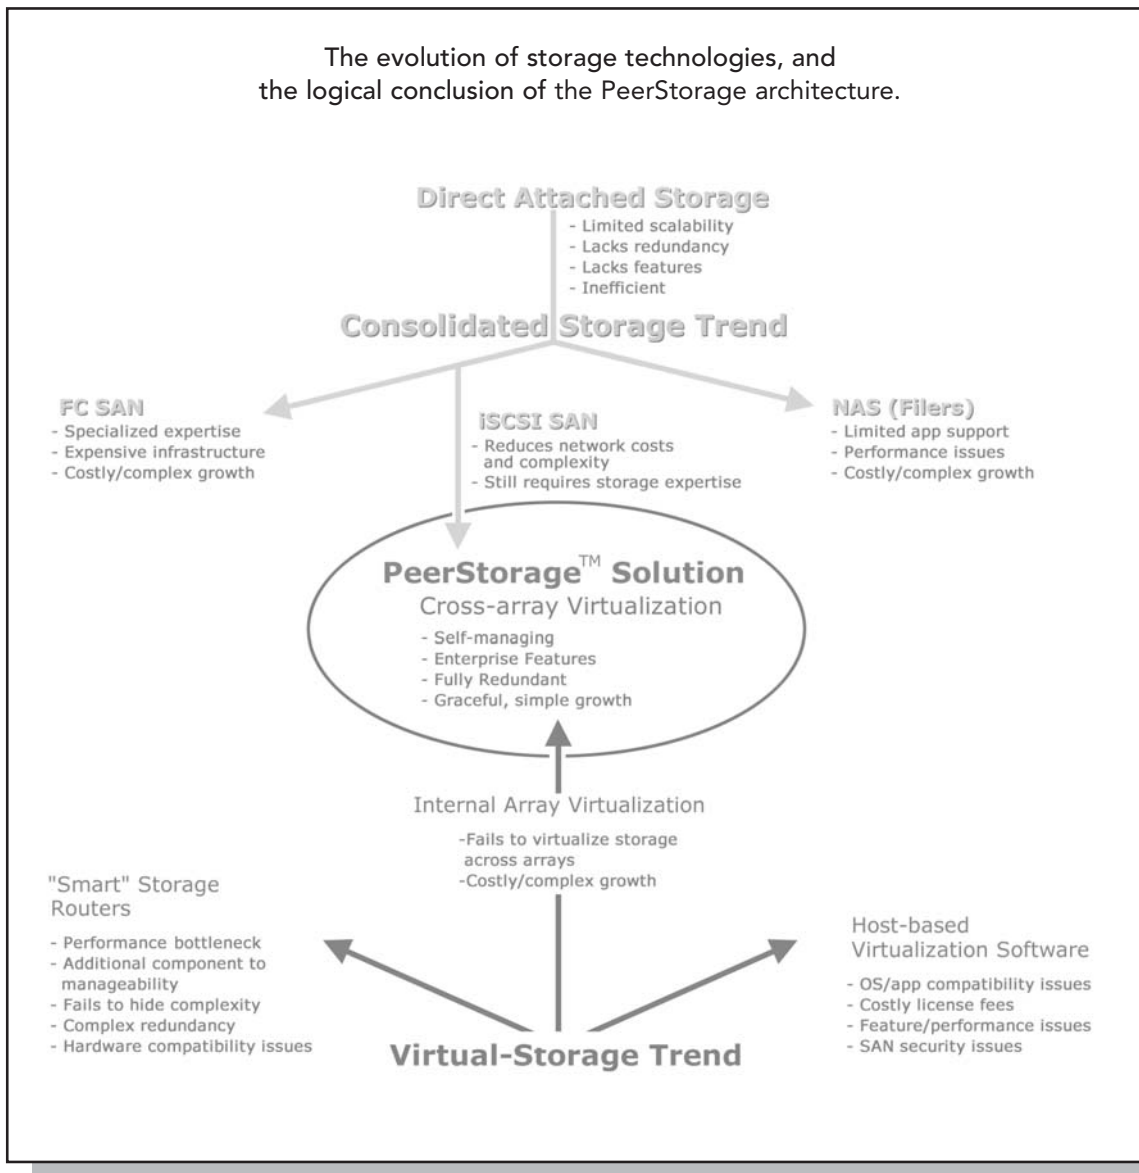
Ultimately, the challenges lie with the storage arrays. But storage arrays cannot manage themselves, and cannot work together to solve the problems.

Or can they?

PeerStorage™ Solution: Self-Managing, Self-Scaling, Self-Healing
PeerStorage 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 PeerStorage arrays transparently to the administrator - and all are accomplished while the system is online, through a common intelligent interface. In addition, PeerStorage architecture enables almost limitless scaling - to well over 100TB of data, significantly more than any SAN or NAS solution can offer. PeerStorage technology uses iSCSI connectivity, delivering that standard's performance and cost advantages.

Rather than attempting to patch archaic storage architectures with bolt-on solutions, the PeerStorage approach addresses the underlying problem: that is, how to manage explosive data growth and meet escalating service expectations, while cutting costs and reducing complexity.



PeerStorage 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.

PeerStorage arrays also automate key functions such as RAID configuration and performance tuning, freeing administrators to attend to other tasks. By exporting a native iSCSI interface, the arrays leverage the simplicity and low cost of Ethernet network infrastructure. Together, these technologies offer a simple, easy to use SAN, without sacrificing enterprise features such as availability, scalability, redundancy, even replication and snapshots.

How PeerStorage Arrays Manage Growth

The proliferation of online data seems to be without bounds; regulatory and legal requirements for data retention compound this problem. Businesses must manage the continual growth of information, growth that is often unpredictable.

The PeerStorage system addresses this problem by automating all aspects of storage expansion. Arrays can be added to existing deployments without disruption of storage service, 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.

Rather than requiring the administrator to parcel out the additional resources, the PeerStorage system automatically qualifies existing workloads and allocates the newly deployed resources accordingly.

This automatic load balancing is a dynamic process, constantly self-tuning the storage system based on ever changing workloads.

Volumes will be automatically distributed across available RAID sets, disks, and disk arrays. Network connections will be automatically balanced across available network ports.

Arrays can be added to existing deployments without disruption of storage service, and the new storage resources are immediately available.

This load balancing is accomplished without any special host software or reconfiguration of host initiators. The distribution occurs entirely at the array level, with no

special networking switches required in the middle. As a consequence of embedding intelligence into the storage arrays, the administrator will be presented with a simple pool of available storage, from which he or she can "point and click" to create volumes and expand existing volumes on the fly.

How PeerStorage 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. Email is a classic example of changing priorities, as preserving email 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. PeerStorage systems 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, and PeerStorage solutions include dual controllers, dual fan trays, dual power supplies, battery-backed cache,

redundant network ports, and multiple I/O network paths. These are critical requirements for any serious storage deployment.

Businesses must manage the continual growth of information, growth that is often unpredictable.

However, the hazards of faulty hardware

present just one of the risks that IT departments face in today's world of hackers, viruses, and server security patch upgrades. Software vulnerability is a tangible danger. Should a software virus strike, reverting to tape backups means hours, if not days, of downtime - something few business can afford. PeerStorage 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 deltas are stored. The result is that at any time, the volume can be "rolled back" to a snapshot; the base volume is restored to the exact state of the volume when the snapshot was taken (e.g., before the virus hit, or the hacker intruded, or the database migration went wrong). This rollback is an immediate operation, taking minutes rather than the hours that tape restoration would require.

No level of software backups or hardware fault-tolerance can protect against the calamity of natural and unnatural disasters. More and more, businesses of all sizes

are exploring disaster tolerance solutions for their data. PeerStorage arrays can remotely mirror volumes at significant physical distances through volume replication across IP networks; this is the only real protection from such events.

From high availability to snapshots to volume replication, administrators need these new tools to provide today's expected services. These features are not only mandatory, they must be easy to implement and maintain. With PeerStorage technology, IT departments can offer these features without adding complexity or cost.

How PeerStorage Architecture Reduces Complexity

An enterprise storage solution that requires constant attention from a team of on-staff experts only serves the largest datacenters. In contrast, most organizations need innovative solutions that hide the technology and simplify operations.

PeerStorage systems do just that: built-in management functionality provides ease of use, and iSCSI connectivity delivers network simplicity with superior performance. Extending this architectural philosophy highlights additional operational benefits:

“Buy-as-you-go” Scalability:

Because PeerStorage solutions are truly modular storage, IT Departments can buy only the storage resources they need today, instead of over stocking disk capacity to avoid the dreaded upgrade scenario.

Robust Features Included:

With management and replication functions built directly into the array, administrators are released from the burden of evaluating complex option packages and compatibility charts just to implement new functionality.

In contrast, most organizations need innovative solutions that hide the technology and simplify operations.

No Host Software Required:

Because of the innate intelligence of PeerStorage arrays, no proprietary host software is needed to access and organize storage, and no dedicated management workstations or specialized network switches are necessary.

Leverage Existing Staff Skillset:

Finally, this network technology lets users leverage current staff expertise in Ethernet IP, while the storage manages itself.

PeerStorage Solutions Deliver Bottom Line Savings

Most important, new storage solutions cannot break the budget. Costs must be kept low and ultimately be justified with a rapid return on investment. The PeerStorage system caps costs while maximizing value.

Low Capital Equipment Costs:

PeerStorage arrays offer the most affordable, fully redundant, enterprise storage array on the market. The low costs of Gigabit Ethernet switches and cabling combine to create the most accessible enterprise SAN available.

No Software License Fees:

Since no host-based software is configured, adding a new server to the SAN does not result in additional software license fees. Also, with high-level software features included in the standard package, there are no "surprise" costs.

Lowest Operational Expenses:

A self-managing system such as the PeerStorage array does not require continuous administrative "care and feeding." The arrays can be setup and serving storage in less than 20 minutes, saving hours and days associated with implementing competing systems.

By reducing expenses throughout the lifetime of the system, the PeerStorage approach dramatically reduces net expenses. Plus, the increased levels of service offered by PeerStorage arrays deliver both lower Total Cost of Ownership and faster Return-On-Investment.

PeerStorage Solution: Finally, a Complete Solution

The next evolution in storage design borrows from past storage technology developments and builds on the innovations of storage over IP. The PeerStorage system provides an intelligent, full-featured storage solution that simply and easily answers the original challenges that launched the storage networking revolution: managing explosive data growth and meeting escalating service requirements without adding complexity or increasing cost.

For more information about PeerStorage solutions or to learn more about EqualLogic, Inc. visit www.equallogic.com



9 Townsend West, Nashua NH 03063
603.579.9762
fax 603.579.6910
www.equallogic.com

Copyright © 2004 EqualLogic, Inc. All rights reserved. EqualLogic is a registered trademark of EqualLogic, Inc. PeerStorage, and Simplifying Networked Storage are trademarks of EqualLogic, Inc. All other brand or product names mentioned are the trademarks or registered trademarks owned by their respective companies or organizations.