



PS SERIES ARCHITECTURE WHITEPAPER

FUNDAMENTALLY ALTERING THE ECONOMICS OF SHARED STORAGE

MANAGEMENT SERIES

Consolidated storage promises many benefits: virtualization, reliability, high availability, and centralized management. Although these promises have yielded interesting technologies, they do not meet the need today for affordable, easy-to-use products and are impractical for most environments.

To solve the business problems associated with storage configuration, scaling, and management, the EqualLogic PS Series was architected to provide a cost-effective, self-managing alternative to expensive and complex offerings – bringing the benefits of consolidated storage to a broad range of markets.

The PS Series Architecture fundamentally alters the economic and management landscape for consolidated storage by combining proven operating system, storage, and networking technologies with patented, intelligent automation for a true state-of-the-art storage system.

This blend of technologies makes consolidated storage a cost-effective reality for all operating systems and applications. The PS Series Architecture uses page-based volume management between arrays and incorporates high-speed, standards-based technologies such as Gigabit Ethernet, iSCSI, Serial Attached SCSI, and Serial ATA. The result is an affordable, scalable storage solution that grows easily from a couple of terabytes to more than 100 terabytes, while management stays minimal and constant.

While the industry talks about “next-generation storage,” EqualLogic PS Series is already there.

The PS Series Architecture was designed to meet several goals:

- Provide a single, shared storage service to hosts and applications
- Scale seamlessly from small to very large configurations
- Automate repetitive management tasks:
 - Perform RAID configuration and sparing
 - Provision new resources in the most efficient manner
 - Allocate resources to meet changing workload demands
- Comply with industry standards
- Require no special host software for access and management
- Deliver uncompromised availability and serviceability

- Offer a full set of features, including capabilities that are usually available only as expensive add-ons, like snapshot and volume replication capabilities
- Provide a cost-effective, easy-to-use replacement for direct-attached storage (DAS)

This whitepaper gives a brief overview of the technologies that make the PS Series Architecture uniquely able to address the issues associated with consolidated storage.

PS SERIES ARCHITECTURE FUNDAMENTALS

To achieve the goal of providing an affordable, self-managing consolidated storage solution, the PS Series Architecture employs three essential elements: the PS Series array, the PS Series group, and volumes.

PS Series Array

The foundation of the PS Series Architecture is a PS Series array – a high-performance block storage device that is ideal for departmental or enterprise storage. Each array is a commonly configured, fully-redundant storage unit containing disks, multiple high-performance network interfaces, controllers with mirrored battery-backed caches, and other advanced features. The disks in an array are automatically protected with RAID (RAID 10, RAID 5, or RAID 50) and hot spares.

PS Series Group

A PS Series group is comprised of a single PS Series array or multiple arrays working together. The group is seen as a single, shared storage service by application servers and administrators. When an array is configured as a group member, its RAID-protected disk space is added to the group’s storage pool. Data and network I/O to the group is load balanced across the group members’ resources.

Volumes

Administrators create volumes from the available space in the PS Series group storage pool. A volume can be spread across multiple disks and multiple group members – this is

done automatically by the virtualization built into the arrays. The group exports volumes as iSCSI targets protected with security, including authentication and authorization, for both discovery and access. Upon connection, hosts see volumes as local disks. Volume snapshots and volume replicas can also be created from the storage pool.

PS SERIES GROUPS FACILITATE STORAGE MANAGEMENT

Although comprised of one or more PS Series arrays, a PS Series Group appears to client servers as a single entity that offers network storage access in block mode – a storage area network (SAN). Unlike conventional SAN products, each group member “cooperates” with other members to automate resource provisioning and performance optimization. This cooperation is a key element in the virtualization of not only disk storage, but also controllers, caches, network links and management of the PS Series group.

The disk data served by a group is dispersed across the group members, with the placement continually adjusted for optimal performance as resources are added or load changes occur. In addition, application server network sessions are distributed across members to spread out the I/O requests.

The result is a storage system that scales linearly in all significant dimensions: storage capacity, network bandwidth, application server request processing capacity, and performance. In contrast, administration remains stable as the group is expanded. There is no separate management of arrays and volumes; all administration is performed at the group level.

With its ability to export many targets, a PS Series group can be used to consolidate storage for a wide range of applications and servers, including e-mail, Web servers, and database servers, and file/print without compatibility or performance concerns and without a prohibitive increase in administrative expenses.

DISTRIBUTED PAGE-BASED VOLUME MANAGEMENT

The RAID-protected disk space that each member contributes to the PS Series group’s virtualized storage pool is partitioned into fixed-sized chunks of data, called pages. Each volume has a page map that allocates and re-allocates pages to the members. A page map is shared by each member that contains a page of the volume. Page maps are also used in the implementation of features like volume snapshots, replication, and over-subscription (thin provisioning).

Page movement is a transactional operation between group members that helps ensure data integrity. Pages can be moved transparently (that is, online and without blocking host access) across disks in a member and across group members for automatic load balancing and to help ensure data availability. For example, to eliminate a frequently-accessed data “hot spot,” the affected pages may be spread across different disks or group members, without the need to move all the pages for an entire volume. The built-in load balancing features use real-time and historical access trends, as well as capacity and status information, to help continuously improve performance within individual arrays and the group as a whole.

A lookup and transfer mechanism is used to efficiently track and process pages located on the group members. When a PS Series group receives a client server request, it identifies the location of the data and transfers the request to the member that contains the data. A request may be transferred one more time if the desired data is located on another member.

The entire page lookup and transfer operation increases latency insignificantly relative to the disk access time. In addition, the PS Series Architecture’s optimization technology continually adjusts storage placement and network connections to optimize performance.

FAST, SIMPLE PROVISIONING AND SCALING

Using the scalability model from the PS Series Architecture, a PS Series group can scale linearly in both capacity and performance – all while the system remains online. The model allows for multiple members in a single group, with a controlled and efficient addition of resources as new members are added.

Individually, each member is a fully-functional, high-performance, highly-available storage array with mirrored write-back caches and multiple storage network connections. Resources like disks, controllers, caches, and network connections can be easily added to and removed from a group with no complex administration tasks or impact on availability.

As capacity and performance requirements increase, a group can be expanded. New members “learn” configuration and performance information from the group – with no manual administrative intervention. Data and application server connection load balancing occur automatically as the group scales. Data access activity is monitored, and data and network connections are adjusted as needed automatically and online.

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 capacity can be added so easily, IT managers need to buy only the storage necessary for today’s applications, easing budget constraints caused by excessive and superfluous purchases. Upgrades to new technologies are done easily as different types of PS Series arrays can interoperate in the same PS Series group at the same time.

Table 1 shows how the scalability model of the PS Series Architecture compares to traditional DAS and SAN solutions.

TABLE 1: STORAGE EXPANSION COMPARISON

EXPANSION ACTIVITY	TRADITIONAL DAS AND SAN	PS SERIES ARCHITECTURE
Add storage capacity or network bandwidth.	Limited expansion capabilities are available.	Add another member to a group; data remains available.
Increase network connections.	Requires repetitive manual tasks.	Minimal administrative effort is required.
Add controllers.	Storage must be taken offline; data is not available.	Expansion occurs online; data remains available.
Perform load balancing.	Requires repetitive manual tasks; data may not be available.	Operation is fully automated; data remains available.

INCREASED AVAILABILITY

Consolidated storage increases the need for availability in the storage infrastructure. The availability model used by the PS Series Architecture essentially eliminates single points of failure and enables an array to survive multiple, simultaneous failures:

- Each PS Series array is composed of redundant components – disks, controllers with mirrored write-back caches, network interfaces, power supplies, and cooling fans.
- Hot swappable hardware components help reduce downtime.
- Faults are well isolated – loss of controllers does not cause loss of RAID protection. This works for all supported RAID types.
- Disks are automatically configured with RAID, and hot spares are reserved to help ensure that no data goes unprotected.
- Hot sparing and data recovery require no user intervention.
- Volume replication can provide site-level disaster protection.
- Group member hardware upgrades can be done while a group is online; simply add the upgraded array to the group.
- A member can be decommissioned while a group is online; simply remove the array from the group.

In addition, when upgrading to the latest technology, IT managers often find themselves writing off relatively new equipment because it is not compatible with new acquisitions. The PS Series Architecture eliminates this problem by ensuring that previously-purchased hardware and advanced technologies will remain interoperable.

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 occur as resources and performance metrics change. No longer must administrators manually map application data to specific physical devices, controllers and network ports.

For example, to create a volume from a group's storage pool, an administrator supplies only the name and size. The group members handle all details of storage allocation and load balancing.

Because a group is seen as a single entity, an administrator has a centralized view of the storage; management tasks remain constant, regardless of scale. A group can be managed through several built-in mechanisms, including SNMP, serial line, telnet, and Web-based user interfaces. No external management station or management software is required, essentially eliminating the need for additional purchases and administrator training.

In addition, a group can alert responsible individuals to management activity or problems through log files, SNMP traps, and e-mail notification methods. Finally, data and management access is protected with authorization and authentication mechanisms.

ADVANCED FEATURES

With traditional SANs, many advanced features are typically expensive add-ons. For example, conventional SAN solutions require server licenses or additional fees for snapshot and replication capabilities. Solutions developed

with the PS Series Architecture include these features at no extra charge.

In addition, the PS Series Architecture's snapshot capability exceeds the snapshot features available on the market today. The real-time synchronization of snapshots across multiple volumes occurs without interrupting applications. Snapshots allow administrators to perform online backups and can be scheduled at regular time intervals. If data loss occurs, archived information can be rapidly retrieved to restore data and return to normal operations.

The storage replication capabilities provided with the PS Series Architecture also enable administrators to replicate storage to a separate location and enhance disaster recovery procedures.

Other advanced features that are included with each PS Series array – without additional license fees or keys – include Auto-Snapshot Manager for Windows® with Smart Copy, multi-path I/O (MPIO), pooled and tiered storage, roles-based administration, storage virtualization, and thin provisioning. In addition, as new features are added to the PS Series over time, customers under warranty or under a service plan are entitled to new firmware releases free of additional charges.

LEVERAGES INDUSTRY STANDARDS

Each PS Series array uses standard Ethernet networks and industry-standard iSCSI protocols. An array has multiple Gigabit Ethernet network interfaces, and each interface performs at 124 MB/second, full duplex. This design lowers purchase and connection costs and – unlike Fiber Channel solutions – greatly expands the number of administrators familiar with the technology that forms the basis of the PS Series Architecture, thus keeping the training required to a minimum.

Any Gigabit Ethernet switch can be used with an array to support most existing storage network infrastructures. By placing the intelligence in the storage arrays instead of the

switch, management of the SAN and storage and network load balancing occurs in an efficient and cost-effective manner.

In addition, an array includes either Serial ATA (SATA) or Serial Attached SCSI (SAS) disks – the latest generation of disk drives – which provide enhanced performance, capacity, connectivity, and reliability. Each drive has a dedicated 300MB/second channel into the array. The drives are hot swappable and come in a variety of models optimized for performance and capacity. The advantages of Gigabit Ethernet, Serial Attached SCSI, and Serial ATA are many, and their feature sets and performance are expected to continue to grow.

The PS Series Architecture also can be used with most leading operating systems, applications, and management frameworks. Any standards-compliant iSCSI initiator can access a PS Series Group by using software implementations included with major operating systems or by using dedicated iSCSI host bus adapters provided by leading storage vendors. Thus, a solution is available for any environment and for every cost or performance need.

Table 2 describes how solutions based on the PS Series Architecture compare to traditional DAS and SAN solutions.

ALL INCLUSIVE BUSINESS MODEL

Another unique characteristic of PS Series storage arrays is the all-inclusive features. A key reason for consolidating storage is to make storage management easier. Other SAN vendors separately license features – working from the assumption that SAN management capabilities are somehow optional. PS Series arrays provide all features standard – this allows for IT staff to make improvements in storage management, backup, data protection as needed, reducing the complexity associated with “out of cycle”

purchases and to reduce delays in projects resulting from storage equipment near the end of its depreciation cycle.

CONCLUSION – SELF-MANAGING, AFFORDABLE STORAGE

By allowing arrays to work collaboratively in PS Series groups and by automating time-consuming manual tasks, PS Series Architecture offers a fundamental change in SAN implementation and operation. The highly virtualized PS Series Architecture is designed to significantly lower the cost of consolidated storage without a loss of capabilities. The result is a full-featured, self-managing SAN that includes dynamic load balancing, hot swapping, hardware redundancy, snapshots, and volume replication.

Storage acquisitions are driven by current application needs, not a prediction of future needs. The PS Series Architecture represents a dramatic advancement in storage economics, from purchase and set up through operation and upgrades. By eliminating the barriers to adopting consolidated storage – that is, non-standard and expensive hardware, complex technology, and high staff skill requirements – the PS Series Architecture provides storage solutions that can quickly grow from terabytes to hundreds of terabytes, using existing resources and familiar skills.

ABOUT EQUALLOGIC PS SERIES ARRAYS BY DELL

The PS Series Architecture is a breakthrough in the process of buying, managing, and growing SANs. PS Series arrays provide affordable, self-managing storage solutions that use iSCSI, Serial ATA, and Serial Attached SCSI technologies with patented PS Series intelligent automation that helps to dramatically reduce storage complexity and operating costs for businesses in need of consolidated storage.

TABLE 2: PS SERIES ARCHITECTURE SOLVES BUSINESS PROBLEMS FACING IT MANAGERS

	DIRECT-ATTACHED STORAGE (DAS)	TRADITIONAL STORAGE AREA NETWORKS (SAN)	EQUALLOGIC PS SERIES SOLUTION
BUDGET	Must plan for both initial and future growth.	Must plan for both initial and future growth.	Buy only what you need now. Can easily expand at any time.
	Low acquisition and management cost (for a single server).	High acquisition and management (people and tools) costs.	Low acquisition and management costs.
	Scaling results in repeated purchases and higher management costs.	Scaling results in expensive infrastructure upgrades.	Low scaling costs.
RESOURCES	Each storage component separately deployed; not a shared storage solution	Shared storage solution; separate storage network infrastructure and tools.	Shared storage solution; utilizes familiar network infrastructure and easy-to-use tools.
	Dedicated.	Many interoperability issues.	Complete Ethernet networking interoperability.
	Upgrades and maintenance cause server and application downtime.	Upgrades and maintenance cause server and application downtime.	Deploy storage on demand and easily change configuration online; upgrade and perform maintenance online.
	Poor utilization of capacity and performance.	Good capacity utilization and performance.	Automated capacity utilization and performance optimization.
PEOPLE	Each storage element is managed individually.	Managed as single storage system.	Each storage element is managed individually on a network, regardless of scale.
	Repetitive, manual management tasks in dissimilar locations.	Repetitive, manual management tasks.	Automation dramatically reduces management time and costs.
	No specialized knowledge and experience required.	Highly specialized knowledge and experience required.	No specialized knowledge and experience required.
TIME	Management effort grows as number of servers grows.	Management effort grows slowly as number of servers grows.	Management effort remains the same as number of servers grows.
	Short training time.	Long initial training time.	Short training time.



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