



PILLAR AXIOM

An Oracle White Paper
September 2011

Pillar Axiom 600: System Overview and Architecture

Introduction.....	1
Benefits.....	2
Highest Utilization.....	2
Investment Protection.....	3
Ease of Use.....	3
Pillar Axiom 600 Hardware Components.....	3
Pillar Axiom Slammers.....	4
Pillar Axiom Bricks.....	4
Pillar Axiom Pilot.....	5
Uniquely Differentiated Technology.....	5
Quality of Service Policies.....	5
Unified Storage.....	6
Distributed RAID.....	7
Policy-Based Management & Application-Aware Storage Profiles..	8
Modular Architecture.....	8
Conclusion.....	9

Introduction

IT executives need to leverage applications to give themselves a competitive edge – and do so in a way that dynamically matches changing business priorities. Because faster drives and faster processors do not automatically equal a higher Quality of Service (QoS) for storage resources supporting business-critical applications, a different solution is needed.

Oracle's Pillar Axiom 600 storage system delivers unmatched capacity utilization, ease of use and an integrated dynamic management of I/O prioritization to make sure critical applications perform optimally under any load condition. It is a modular storage platform built on three intelligent hardware assemblies: Pillar Axiom Slammers (storage controllers), Pillar Axiom Bricks (drive enclosures), and Pillar Axiom Pilot (management platform). The Pillar Axiom 600 storage system with QoS technology allows you to prioritize data access based on the significance of key applications to your business. With Pillar Axiom's patented QoS technology, IT executives like you can complete more key business initiatives with fewer storage headaches.

Benefits

Pillar Axiom 600 delivers performance, availability, protection, capacity, and scalability for multiple tiers of data. Virtualization of the storage pool allows effortless expansion and management of the system, while the unique Quality of Service (QoS) allows Pillar Axiom 600 to direct resources to applications most critical to your business. The Axiom accomplishes this by controlling not only I/O priority but managing all aspects of the entire I/O stack. This approach to storage dramatically simplifies data management and offers a price to performance ratio unparalleled in the industry. Additionally, QoS settings may be saved as Application Profiles and reused to quickly and optimally provision storage for new or additional instances of any given application.

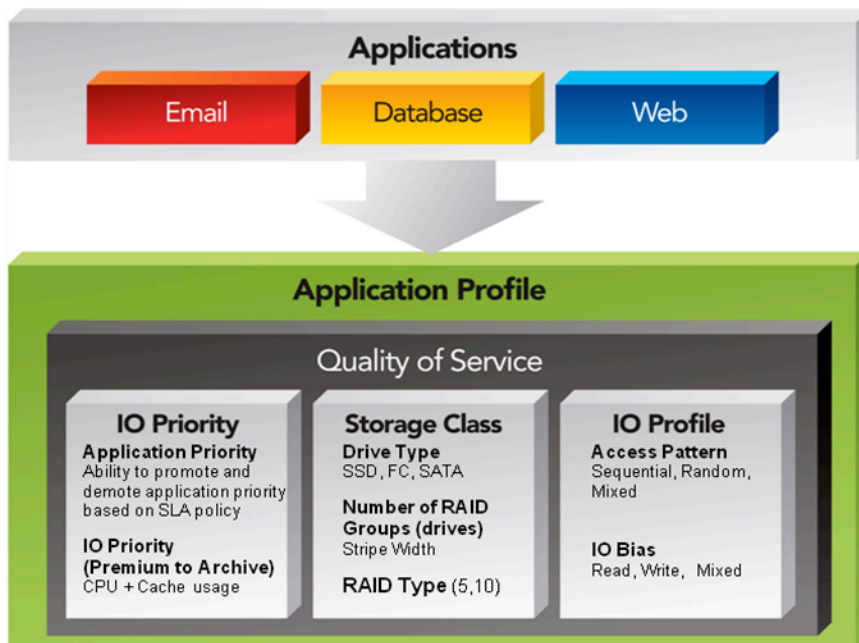


Figure 1. The combination of I/O Prioritization, multiple storage classes, I/O profiles customized to specific application needs, and ease-of-use results in the delivery of application-aware deterministic performance under any load condition.

Highest Utilization

Traditional storage systems are not designed with application importance and priority in mind. Hence, they require applications of different priorities to be segregated on different drives drastically reducing capacity utilization to ensure performance requirements are met.

Storage administrators have little control over where data may physically reside on disk. With increasingly chaotic access patterns, “hot spots” are neither predictable nor controllable with most all legacy architectures. To insure that storage delivers consistent performance, the administrator is often forced to limit the overall workload to much less than the theoretical maximum to insure that there is a very low chance of any given storage system becoming over saturated and causing I/O bottlenecks.

Pillar Axiom 600 incorporates QoS with application I/O prioritization that enables consolidation of multiple tiers of applications on to a single storage system. By prioritizing data access and dynamically managing any I/O contention, capacity can be used more efficiently, thus increasing storage utilization without sacrificing performance.

Investment Protection

Business requirements and priorities keep evolving constantly. Often these changes necessitate a rip-and-replace upgrade. With Pillar Axiom 600, customers can build a storage infrastructure that has the ability to scale both capacity and performance quickly. Pillar Axiom 600 can adapt—in a non-disruptive way—to the ever changing needs of the business by

- Incorporating latest advances in storage controllers by allowing legacy and modern technologies to coexist within the same platform
- Seamlessly adapting to changing priorities of application by tuning QoS levels on-the-fly
- Storing application data on different tiers of storage and accessing the data stores through different protocols, thus allowing storage consolidation and eliminating storage islands.
- Providing the ability to easily move application data off of older technology media to new media and allow the removal of the old media; all in a non-disruptive manner.
- Providing robust tools for data protection caused by human errors or natural disasters

Ease of Use

With the industry's most intuitive storage management user interface, Pillar Axiom 600 can be deployed, provisioned, managed, and maintained easily without special training. Key technologies that contribute to simplified management include:

- An intuitive, complete graphical user interface (GUI) for management that enables an IT generalist to quickly provision and re-provision data volumes and file systems without requiring specialized storage knowledge.
- Complex, mathematical functions integrated into the Axiom that select optimal data placement, control cache operations, CPU cycle allocation and I/O ordering thus eliminating the need to manually tune and optimize your storage. If the initial assumption is incorrect or requirements changes, Pillar Axiom 600 can be reset with a couple of mouse clicks without disruption to application data access.
- Built-in guided maintenance allows many maintenance tasks to be performed by staff without requiring a trained technician making an onsite visit.

Pillar Axiom 600 Hardware Components

The Pillar Axiom 600 is a modular storage platform built on three intelligent hardware assemblies:

- Pillar Axiom Slammers (storage controllers)

- Pillar Axiom Bricks (drive enclosures)
- Pillar Axiom Pilot (management platform)




Figure 2: Pillar Axiom 600 building blocks

Pillar Axiom Slammers and Bricks communicate through the Storage System Fabric (SSF), a switched fibre channel back end mesh, which allows any Pillar Axiom Slammer controller to access any Pillar Axiom Brick, providing a shared storage pool across all storage controllers.

Pillar Axiom Slammers and Bricks can be flexibly combined to meet unique application performance and storage capacity requirements. Advanced quality-of-service software manages system resources (CPU, cache, and capacity) to automate storage provisioning based on business priority.

Pillar Axiom Slammers

Pillar Axiom dual redundant Slammers include the following components:

- Asymmetric Active-active SAN or NAS control units
- 48GB of cache
- Eight GigE ports or four 10GigE ports (NAS)
- Four 8 Gbps FC ports or four 4Gbps FC ports and/or 4 or 8 GigE iSCSI ports (SAN)
- Complete high availability
- Scales up to four Pillar Axiom Slammers (eight control units)

Pillar Axiom Bricks

A Pillar Axiom Brick is a fully redundant storage enclosure, used as a modular building block to create a shared storage pool which the NAS and SAN controllers in the Pillar Axiom Slammers can access through the SSF. The Pillar Axiom 600 supports SSD, FC, and SATA drives. Each Pillar Axiom Brick provides health and performance monitoring and diagnostics.

Within a Pillar Axiom SATA Brick, the SATA drives have a unique Pillar designed multiplexing mechanism that allows it to be accessed by both RAID controllers. The Pillar Axiom SATA Brick and Pillar Axiom SSD Brick contain 13 dual-ported drives, two RAID controllers and dual redundant power & cooling modules. The drives within a Pillar Axiom SATA Brick are configured as two 6-drive

RAID groups which support RAID 5, 50 or 51 and RAID 10, with the thirteenth drive used as a hot spare.

The Pillar Axiom FC Brick contains 12 dual-ported drives configured as an 10+1 drive RAID 5 group, or an 11 drive RAID 10 group with the twelfth drive used as a hot spare.

Pillar has placed a hot spare in every brick to localize the impact of RAID rebuilds to only the effected RAID group so as to not impact any other RAID groups. This allows for a vastly reduced performance impact during rebuilds and much faster rebuilds; the fastest in the industry.

Pillar Axiom Pilot

Pillar Axiom Pilot controls all system configuration, management, and monitoring. The Pillar Axiom Pilot uses an Ethernet-based, out-of-band private management interface (PMI) to communicate with the Pillar Axiom Slammers. Access to the Pillar Axiom Pilot is provided through a GUI or CLI. The Pilot consists of two independent control units (CUs) that operate in an active/standby mode. A heartbeat runs between each CU. In case one CU fails, the other CU becomes active and takes over Pillar Axiom 600 system management.

Uniquely Differentiated Technology

Pillar Axiom 600 allows you to optimize performance based on application I/O needs. It is the first and only storage array that does not use simple FIFO I/O queuing that treats all I/Os on a first come first served basis regardless of the value of the application issuing the I/O. QoS enables intelligent sharing of storage resources, CPU, capacity, and cache so your applications can deliver optimum deterministic performance in multi-tenanted environments. Through innovative QoS software, you are also able to differentiate between drive types in the storage pool to increase application performance and/or reduce costs. This section discusses the unique technologies that differentiate Pillar Axiom 600.

Quality of Service Policies Based on Dynamic CPU, Cache, and Capacity Management

The fundamental design point of the Pillar Axiom 600 is to prioritize all of the storage system's resources based on an application's value. I/O prioritization allows the user to specify, by application and LUNs/Filesystems of the applications, the desired QoS level required to meet the business needs. This QoS level policy controls the allocation of all storage system resources in the data path to meet the business value of the application.

The QoS polices allow the user to specify to the system the expected I/O patterns of each application (random, sequential, read- or write-biased, and mixed). Pillar Axiom 600 will dynamically optimize RAID type, the physical data placement on the optimal drive type (SSD, FC, or SATA) and caching mechanism so as to optimally use all system resource in the most efficient manner.

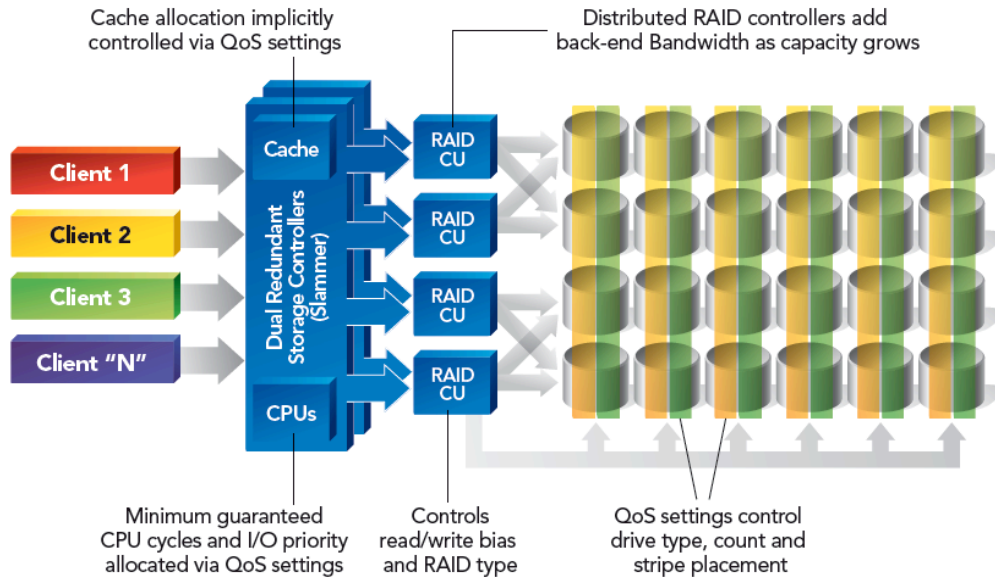


Figure 3. Quality of Service control points

Because workloads and priorities often change, all QoS policies can be changed at any time to reflect changed business conditions. This can be done either temporarily or permanently. A temporary change will change the CPU, cache, and I/O priority settings, but will not actually move data to a new storage class. A permanent change may physically move the data to a different drive type (if requested) or rebalance the storage extents to increase or decrease the stripe width. This data movement takes place in the background, and the data remains available.

Unified Storage

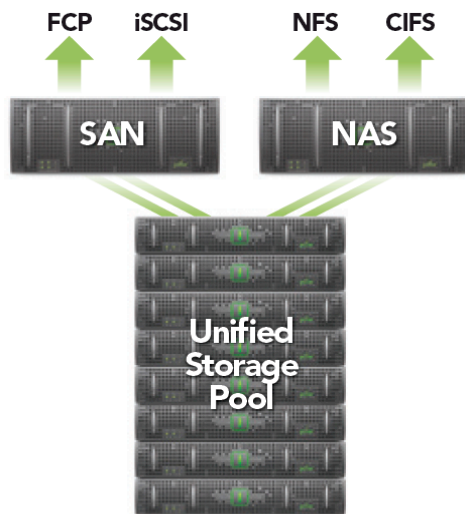


Figure 4. Single Storage Pool for both SAN and NAS storage

For years, enterprise datacenters were forced to store SAN- and NAS-based data in individual silos, an inefficient practice born out of necessity for those who needed to utilize these different protocols to store block and files data. As data volumes grew in these storage area networks (SAN) and network attached storage (NAS) environments, IT managers were required to scale capacity for these different platforms independently and incurred higher capital expenses from having to scale-out two distinct data silos. Due to the specialized expertise required to manage the different types of data and systems, operation costs grew excessively as data centers had to hire two different administrators.

Pillar Axiom 600 is a native, multitenant storage solution not just for SAN or NAS. Customers can choose native NAS or SAN storage controllers—known as Pillar Axiom Slammers—depending whether their applications are file based or block based. The Pillar Axiom 600 storage system uses a single storage pool that can contain Solid State Drives (SSD), Fibre Channel (FC) drives, and Serial Advanced Technology Architecture (SATA) drives. Virtual LUNs that are created in the global storage pool are presented as file systems in a NAS environment and as SAN LUNs in a SAN environment. Storage administrators are now able to consolidate onto one unified storage system and assign a QoS level to the data. QoS, not protocol, is the criteria that determines how the data is stored—the performance and responsiveness that the application requires.

Distributed RAID for Linear Scaling, Performance Under Fault, Fast Rebuild

Storage has evolved from the direct-attached model to a shared-resource model delivered either through an FC SAN or IP-based NAS. SAN and NAS storage continue to grow seemingly without end. These shared-storage pools must now deal with the data of multiple applications converged onto ever-larger multi-tenanted storage systems. Something drastic had to change in the way storage systems have been architected in the past.

The innovative Pillar Axiom 600 architecture places RAID controllers within storage enclosures—called Pillar Axiom Bricks. This distributed RAID architecture provides superior performance compared to legacy approaches that can have as many as 1000 drives assigned to as few as two RAID controllers. The Pillar Axiom system can have 128 individual RAID controllers all operating in parallel. Compared to a system with just two controllers, the performance implications are obvious.

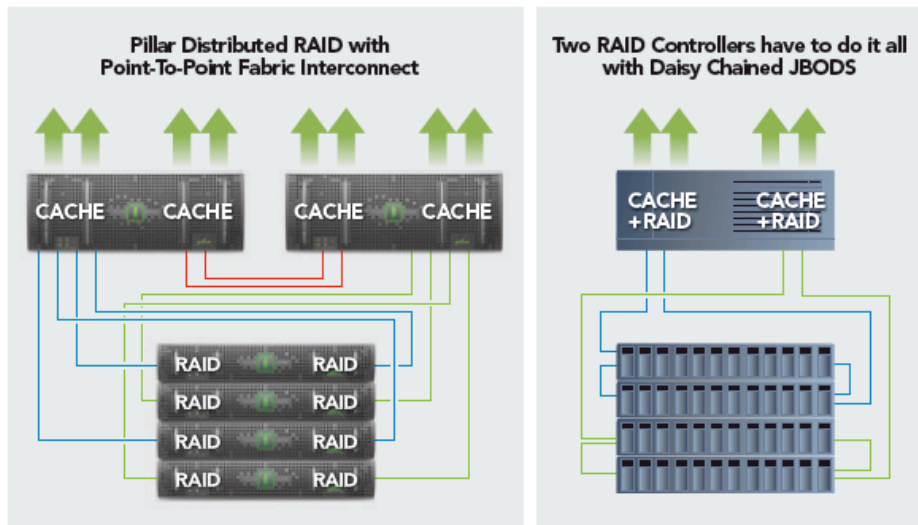


Figure 5. Distributed RAID architecture

Pillar Axiom Brick storage enclosures include built-in dual RAID controllers, removing the RAID functionality from the IO Controller into the storage enclosures. Moving the RAID infrastructure into the storage enclosure reduces the overhead of RAID calculations and dramatically improves drive rebuild times. This ensures that performance scales linearly with increasing capacity. Said differently, as you add capacity, you don't just add JBOD drive trays, you add two dual redundant RAID controllers with every storage controller. A fully configured Axiom can perform 128 concurrent RAID XOR operations vs. just two for most all other modular storage offerings.

Policy-Based Management and Application-Aware Storage Profiles

Pillar Axiom 600 manages system resources to simultaneously deliver multiple levels of service from a single storage pool. The Pillar Axiom QoS manager will automatically configure CPU resources, spindle striping levels, mirroring, priority queues, network bandwidth, and caching algorithms to deliver the appropriate level of application performance based on business requirements. It maximizes the use of the following Axiom system resources: CPU, Cache and Capacity as well as I/O priority thus aligning system performance with application value.

Modular Architecture: Scale with Highest Utilization Rates in Enterprise Storage

Pillar Axiom 600 provides unmatched scalability, scaling from an entry level configuration that has a single Pillar Axiom Slammer and a single Pillar Axiom Brick to a system that includes four Pillar Axiom Slammers and 64 Pillar Axiom Bricks and that can provide up to 1.66PB of storage. Scaling up to four Pillar Axiom Slammers allows the Pillar Axiom 600 to support up to 32 x 1 Gbps Ethernet interfaces or 16 x 10 Gbps Ethernet interfaces for NAS and 16x 8 Gbps FC interfaces or 16x 4Gbps FC interfaces or 32x 1 Gb iSCSI interfaces for SAN configurations.

Previously, if application performance requirements were not being met, an entire new system would be purchased; a procedure commonly referred to as "Rip and Replace". This still holds true for most

competitors' arrays. Aside from the cost savings in equipment and software licenses that the Pillar Axiom 600 provides, this consolidation of applications results in lower power and cooling, management and maintenance costs as well as higher performance. This results in not just capital expenditure savings but also significant operating expenditure savings as well. Customers have the option to increase the throughput of the array independent of capacity, by simply adding more Pillar Axiom Slammers to the system. This is referred to as multi-dimensional scale capability. With true unified storage and unique multi-tenant architecture, customers can chose to add native NAS or native SAN functionality to existing systems sharing the same storage pool and utilizing the same management interface.

Conclusion

Oracle's Pillar Axiom 600 is the most cost-effective, networked storage solution with the highest utilization rates in the market. It is the first truly application-aware system, enabling multiple tiers of storage on a single platform—each with a unique level of service. The result is an ultra high-capacity storage system which consolidates assets, modularly scales capacity, and reduces management complexity.



Pillar Axiom 600: System Architecture and
Overview
September 2011

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2011, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0911

Hardware and Software, Engineered to Work Together