

Pillar Data Systems Defies and Redefines Storage Management Traditions

Analyst: Anne MacFarland

Management Summary

Storage has traditionally given the administrator many things to do. A heritage of tuning and redundancy schemes was developed in a time when capacities were constrained, and disks were slow and subject to frequent failure. Improvements across all these areas has made these traditional tasks more automated and more transparent, and has facilitated the development of other data services such as integration, migration, indexing, etc. Nevertheless, the improvements have not eased the need to optimize the input/output (I/O) of data access. The enterprise need for real-time customer responsiveness demands continuous effort to drive latency out of information systems. And, in the end, assured, fast data access is the basic focus of storage management. Replication utilities and data-services are tools, but matching the resource to ever-changing needs (i.e., tuning) is the key.

With automation and processor speeds what they are, this data access optimization need not be done on a disk-by-disk-by-disk basis. Indeed, it is more efficiently and effectively done more broadly – across an aggregation of disks, even pools of arrays – if you have the algorithms and policies in place to deal with complex optimizations for different priorities of applications and simply “make it so.” Moreover, if you can do that, you can break several traditions that keep storage management a tedious, fragmented, time-consuming, and expensive task.

- **The tradition of buying different storage products** depending on who you are, or on the immediate need. The former, in the past, has kept smaller enterprises from getting the functionality that they needed, and has left larger enterprises with the sour feeling that they probably paid too much for their data storage. The latter requires a prescience that is challenged by the increasingly dynamic nature of many enterprises and workloads.
- **The resulting tradition of a disaggregated management environment** of a hodge-podge of in-array data replication and data movement capabilities and array-spanning point products. Even aggregated management on a single pane of glass does not reduce the vast and complex array of knowledge an administrator needs to make intelligent decisions.
- **The resulting drudgery of configuring arrays on a disk-by-disk basis**, and the difficulty of responding to change in the environment, either at the infrastructure or business levels.

Pillar Data Systems, of San Jose, CA, seeks to enable a break with these traditions by building configuration algorithms and quality of service metric into their *Pillar Axiom* systems. By renaming and refocusing the roles of storage system components onto an I/O engine called a *Slammer* (traditionally known as a controller), an I/O target known as a *Brick*, and a *Pilot* to manage their interaction, Pillar takes a lot of particularities out of the storage management focus, which is moved to a system level. Tuning, the principal task, is done across an entire system. This changes how storage capacity is used and how it is bought. Read on for the details.

IN THIS ISSUE

➤ Storage Management Begs To Be Rethought	2
➤ Pillar Data Systems Provides A New Approach	2
➤ Elements of Axiom.....	2
➤ Pillar’s Approach To The Enterprise Market.....	4
➤ Conclusion	5

Storage Management Begg To Be Rethought

Disks differ by type, brand, and vintage. Assigning storage arrays to tiers based on their capacity and capability makes sense in a time when disk speeds and capacities are rapidly improving, and storage needs are growing. Tiered storage has been used as a coarse alternative to the tuning and data placement that were done to get jobs run when resources were constrained, and time for such tasks was more plentiful. Throwing tiered capacity at a problem of too much data can bring hardware costs under control, but it replaces it with a problem of too many particulars to be managed.

Tuning by tiering is also less simple than it used to be. There was a time, when data use was ruled by the generating application, that the life-cycle of information could be plotted as a downward curve of usage. New retention requirements have changed the predictability of data use – and so has collaboration, partnering, and trying to support additional revenue streams with the same assets. Modeling can explore the what-ifs, but there are many uncontrolled variables in the data generation and usage equation, such as new business strategies, sudden customer demand, or unexpected litigation.

What would be best would be to be able to tune the quality of service profile of an application once, have that profile applied across the entire storage system, but also have the optimization applied at a granular level, using the tuning techniques of striping and data placement, to get the best use out of capacity. To have this approach work efficiently, the environment would have to be engineered to support multiple appli-

cations of various priorities. The naturally slower inner bands could be used for workloads more tolerant of latency, and the faster outer bands could be assigned to the more mission critical applications. A single system could provide different qualities of service (response time, data protection, and other data services) to meet the various need of the enterprise.

Pillar Data Systems Provides A New Approach

Pillar Data Systems' *Axiom* storage solution does just that. Administrators can specify what is needed for each application, using Pillar Quality of Service (QoS) metrics. (See Exhibit 1, in column at left.) Pillar then provides the software to implement the profiles efficiently, providing multiple qualities of service, and flexible qualities of service, in a single array. (See Exhibit 2, on the next page.) The co-centric QoS zones on a spindle are not fixed, and can expand or contract as needed. The ability to monitor and modify data streams on the fly will added to the product in versions 2 and 3, expected later this year.

Elements of Pillar Axiom

Pillar Axiom Slammers, as the controllers are called, are personalized for SAN (block access) or NAS (file access), as needed. Each Slammer connects to each Brick in the system. A Pillar Axiom Brick consists of the ports to support for any-to-any access with the slammers, and two RAID groups plus a common spare drive of SATA storage devices. Pillar Axiom's Pilot manages the Slammers and the Bricks, and defines the extent of an Axiom system. (See Exhibit 3, on page 4.) At first release, a Pilot will manage two Slammers and sixteen Bricks. In the next release this summer, it will manage four Slammers and thirty-two Bricks. Next year, a Pilot will manage four Slammers and sixty-four Bricks.

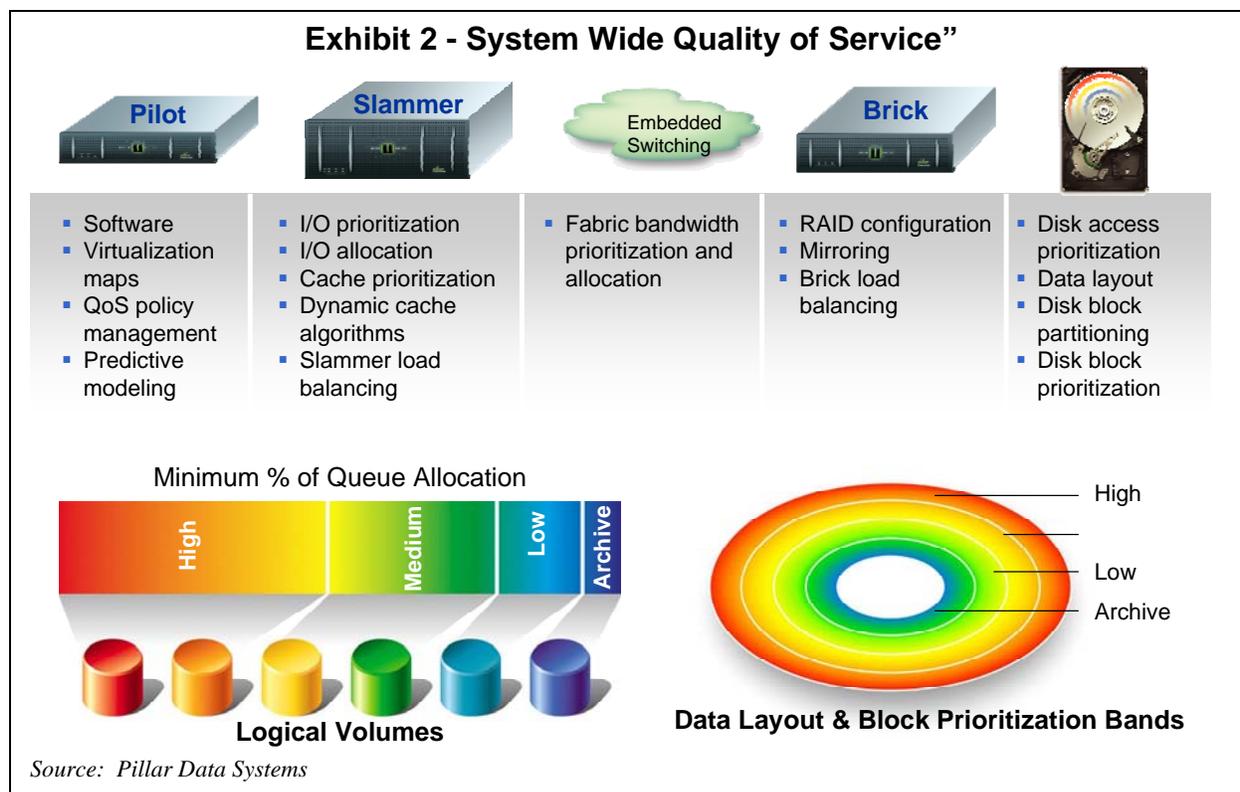
The connectivity is gigabit (gb) Ethernet and 2 gb Fibre Channel. Pillar will look at 4gb Fibre Channel and at ten-gigabit Ethernet in future releases. The Pillar Axiom system offers layers of security (AES, 3DES), snapshot and asynchronous replication, and failover drivers¹ between Slammers as part of their solution. Pillar Axiom has remote call-home capabilities, and the SMTP traps to feed HP *Open View*, IBM *Tivoli*, CA *Unicenter* and *BrightStor*, and other management frameworks. Unlike some vendors, Pillar

Exhibit 1 – Pillar Axiom's Quality of Service Categories

- **Performance Metrics** – The response time needed
- **I/O Bias** – Reads versus writes (which determines cache size)
- **Redundancy requirements** – Default RAID 5 or the additional protection of double or triple redundancy.
- **Scalability** – the rate of data growth
- **Capacity** – LUN or file size

These are known characteristics to application/database administrators.

¹ Pillar claims that at first release, both SAN and NAS restart in under 10 seconds, far faster than the competition.



endowed the Axiom GUI and API with exactly the same complete set of control points, so that third party applications will be able to integrate with Pillar Axiom’s capabilities. Pillar Axiom will support SMI-S by next year, by which time there will be more capabilities included in the specification.

The Pillar Axiom system is architected for multi-tenancy by data from many applications – in fact, that’s what makes it work more efficiently than other arrays. Because for every high-priority workload, there are other workloads, within the same application or in other applications, that are lower priority. This allows well-targeted full utilization of each spindle.

Pillar Axiom is Different

What is different, besides the nomenclature, is how these modules function together and are managed.

The Pilot (Management)

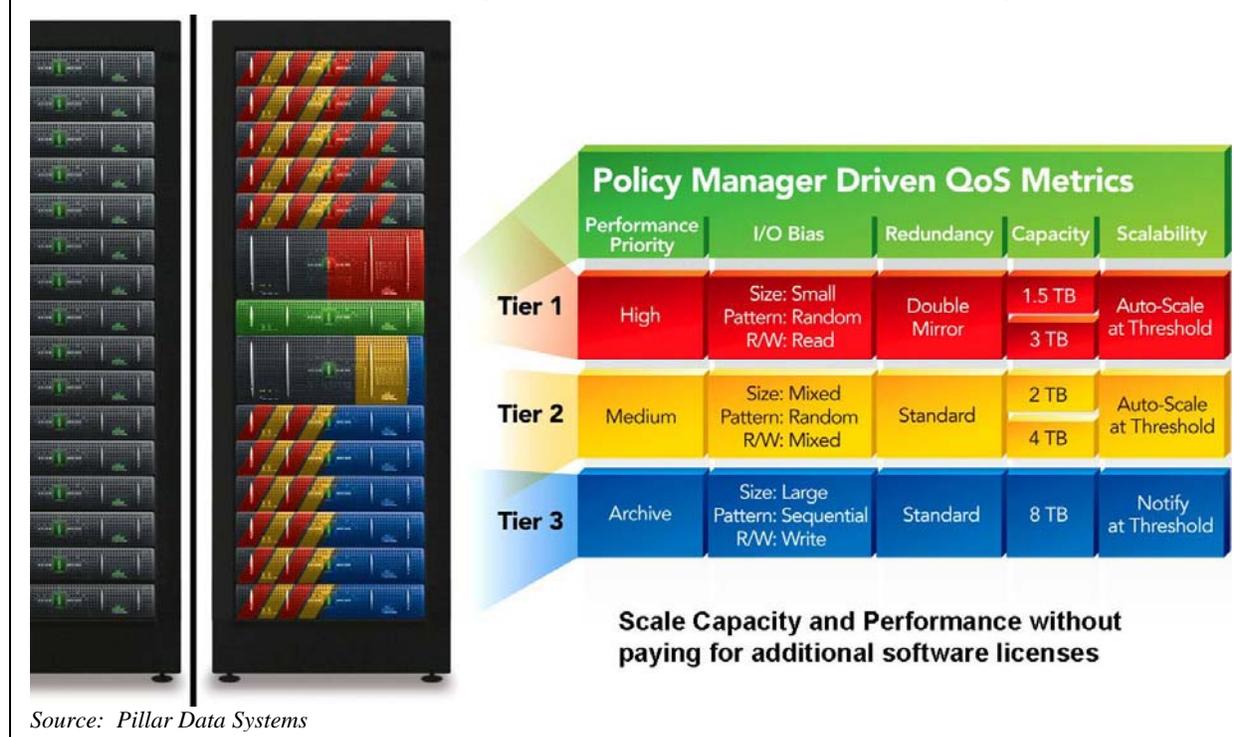
The Pilot (the green box in Exhibit 3) is an active-passive, dual-processor policy engine. Active-passive delivers high availability and easier failover. Storage policy should not be subject to frequent unexpected spikes, so an active-active configuration would be needlessly complex.

The Pilot GUI has been designed for ease of use. When building the Pilot GUI, Pillar Data Systems trained cameras on the face, screen, and keyboards of their system administrator guinea pigs and modified their design to facilitate true ease of use. NAS and SAN have the same screen shots and the same navigation in the Pilot GUI.

- Provisioning mode includes creating file services, creating and deleting LUNS, creating WORM volumes, and predictive modeling to see what effect all this creativity will have on the system. Activity at this screen should be sporadic.
- The Management screen is where there will be daily activity, including oversight, fine tuning, LUN and File System growth monitoring, snapshots and backup scheduling, and fault logs.
- Administrators can beacon metrics of interest for proactive alerts. The monitoring of beacons is done by Pilot, so it is out of band.

At first release, reprioritizing an application, *i.e.*, taking application files (NAS) or LUNS (SAN) and copying it to a higher priority location, is a manually invoked intervention, as is now true on most arrays. In the next software release, it is scheduled to become automated.

Exhibit 3 – A Physical View of Axiom’ Multi-tenancy



The Slammer (Controller)

The Slammers are the 4U boxes in Exhibit 3. Physical space allocations, and most of the virtualization, are done in the Slammer. As Axiom grows to more Slammers and more Bricks, the ability to aggregate more premium space will allow it to be used for the data of demanding applications

One benefit of being able to use disk resources very creatively is that, depending on the QoS desired, SATA can out-perform FC for streaming – by putting premium data on only the best tracks and spreading it widely. (See Exhibit 4, on the next page.)

The Brick (Disk Enclosure)

Two RAID groups and a common spare drive is the basic cellular unit of a Brick. The front-side switch allows any Slammer (SAN or NAS) to access any drive, which increases the opportunities to optimize any applications data. More to the point, the capacity of the Brick is considered a resource, not an assignable asset. It is virtualized by the Slammer. Data doesn’t get an assigned parking space, but it gets the space it merits, with the management to assure that the less deserving cannot poach it.

Like the lesser painters in the atelier of a Renaissance master painter, Pillar Axiom is the

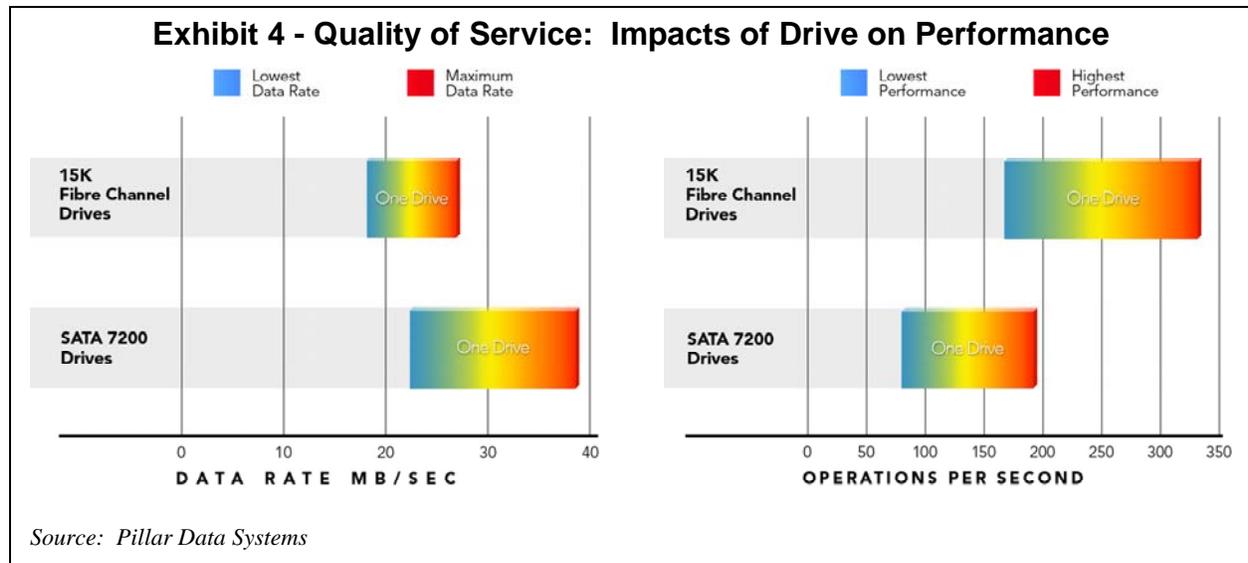
studio drudge that lets administrators craft the impactful aspects of a storage environment – the fast response times and throughput (like the face of a portrait) while taking care of the background of routines and consequences.

Pillar’s Approach To The Enterprise Market

Pillar Data Systems aimed to design a storage system that can satisfy the needs for all tiers of data access service without compromise. The initial release focuses on the lower tiers, for Pillar knows that the second- and third-tier application is where enterprises would want to try out this new approach. They feel confident that once customers see what the Pillar approach can do, they will adopt Pillar’s systems broadly in the enterprise.

Pricing

Pillar Axiom is being offered at 40% below mid-tier prices – just higher than the prices of utility storage. Having thus done an end-run around the rigid customer classification that rules most channel structures, Pillar Data Systems also disrupts the status quo by offering a one-time software license – a one-time charge for the system. This single license spans up to four Slammers – the equivalent of eight stand-alone



NAS systems. The more traditional renewable license, based on terabytes-under-management, has been taking an aggressive bite out of enterprise budgets in times of rapid data growth. The one-time software license may promote enough goodwill to get a trial at data centers daunted by a future of rising software costs.

All hardware elements of the Axiom are field-replaceable. Together with the no-single-point-of-failure architecture, this provides great resilience. The software will be updated periodically by download (remember that one-time license fee), giving the product longevity in the field as well.

Pillar Data Systems provides 7x24 technical support and is partnering with IBM Global Services for extended field support. Each customer has a personalized support web site, and a lot of thought has been put into populating it usefully. This is a low-cost, high-touch approach that, if properly extended to reflect customer experience and preferences, should make the support process more pleasant for all involved.

Pillar plans to use both resellers and direct sales, with a cooperative program that is designed into the business model from the start. The channel should find the Axiom an effective component for all-in-a-box solutions for the smaller enterprise, since Pillar Axiom is a glutton for mixed workloads.

Conclusion

Change is needed in storage management. Pillar Data Systems changes the focus by allowing a more direct optimization of multiple service levels, by unification of SAN and NAS

into a single resource pool, while providing a highly redundant, rugged system.

It can be hard to adopt new ways with old tools, haunted by the ghosts of former habits. A change of vocabulary, and carefully crafted supporting automation, may be what is needed to let storage administrators do the important tuning of the system on a system level, rather than on a disk-by-disk basis. If using a Pilot to craft the way the Slammers will hit the Bricks is appealing, go for it. You will be doing the job right.



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Anne MacFarland is Director of Infrastructure Architectures and Solutions for The Clipper Group. Ms. MacFarland specializes in strategic business solutions offered by enterprise systems, software, and storage vendors, in trends in enterprise systems and networks, and in explaining these trends and the underlying technologies in simple business terms. She joined The Clipper Group after a long career in library systems, business archives, consulting, research, and freelance writing. Ms. MacFarland earned a Bachelor of Arts degree from Cornell University, where she was a College Scholar, and a Masters of Library Science from Southern Connecticut State University.

- ***Reach Anne MacFarland via e-mail at Anne.MacFarland@clipper.com or at 781-235-0085 Ext. 128. (Please dial “128” when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, and “*clipper.com*” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, “*Navigating Information Technology Horizons*”, and “*teraproductivity*” are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.