



Worry-Free Storage — Pillar Optimizes Storage for the SMB Budget

Analysts: David Reine and Anne MacFarland

Management Summary

The typical SMB data center is amazing – well, perhaps we should simply say, it is a maze. It is a maze of under-utilized, heterogeneous, open systems servers deployed throughout the enterprise with a single application installed. It is a maze of under-utilized storage systems from any number of vendors, each clamoring for additional resources. It is a maze of networking devices required to interconnect this complex infrastructure to the SMB's partners and customers alike. It is a maze of cables, contributing to the complexity of the environment, and to the confusion of an over-worked administrative staff. In order to simplify this IT architecture, many data centers have already undertaken the daunting task of changing the paradigm of the data center, consolidating and virtualizing their server sprawl, deploying multiple applications on a single server, each running in its own virtual machine. Consolidation and virtualization of the servers enables the IT staff to improve the utilization of server resources, from under 20% to over 75%, making better use of the expensive compute power and memory resources, which are being wasted every day. Consolidation and virtualization of the server farm also has enabled the data center to reduce the power consumption required to run business- and mission-critical applications and to cool the environment through the use of energy-hungry HVAC equipment, lowering the total cost of ownership (TCO) of the infrastructure. Many SMBs have already recognized this opportunity, implementing these cost-saving measures. However, what is being done to reduce the costs and take the worry out of storage arrays throughout the enterprise?

Unfortunately, the typical storage array is operating at an efficiency level comparable to that of an under-utilized server, around 30%, or less. Storage administrators allocate space on newly installed arrays based upon what they will need, not what they do need. As a result, the majority of the disk storage space remains unused, while other applications may be clamoring for additional storage resources. As with the server farm, this places an unnecessary burden on facilities energy demands. Furthermore, when multiple applications are consolidated on a single server, disk I/O requests do not go down; they go up, placing even more demand on an unprepared SAN. These demands are also heterogeneous, with OLTP applications contending with HPC applications for high performance disk space, while an auditor may be trying to retrieve an archived file from a high capacity device, each carrying a unique cost impact. What is being done to consolidate and virtualize the storage component of the IT equation to ensure that data is assigned to storage on a value basis?

One company that is trying to lower the TCO of the storage network is Pillar Data Systems. An eight-year-old company based in San Jose, CA, Pillar has developed a modular system, *Axiom*, designed to improve the utilization rates for data center storage while lowering the TCO of this expanding resource. To learn more about *Axiom* storage and the *AxiomOne* software suite, please read on.

IN THIS ISSUE

➤ Storage Needs in the Data Center.....	2
➤ Pillar Addresses Three 21st Century Information Challenges.....	2
➤ Pillar Axiom Family.....	3
➤ Pillar AxiomONE Software Suite	5
➤ Conclusion	5

Storage Needs in the Data Center

With the possible exception of applications, the SMB data center now spends more on storage than any other component. Because of the growth of business databases, demands of compliance officers, requirements for business continuance, and expanded customer requests, data center storage has been expanding as never before, with some sites doubling in capacity needs every 18 months, or sooner. Unfortunately, as with the pre-virtualization deployment of open systems servers, the data center has deployed storage badly. Unused capacity has been locked away for some future requirements that may never arise. Under-utilized arrays waste resources of all kinds: disk capacity, energy, floor space, and administrative staffing. At the same time that servers are being consolidated and virtualized¹, improving server utilization, storage utilization rates are suffering, both as a result of the server consolidation and the limited functionality of yesterday's storage.

Server virtualization, in fact, can have a negative effect on storage as the number of I/O requests to the array accelerate and expensive storage resources, such as SSD and FC devices, are consumed by second- or third-tier applications. Applications that previously had dedicated storage, either SAN or NAS, now find themselves in virtual machines on the same physical server, sharing the same storage resource. **Better storage utilization is no longer just a wish; improved utilization of storage capacity is now mandatory! Under-utilized storage can cost the enterprise real dollars on the bottom line.** It can also cost a vulnerable CIO his job.

In order to lower the TCO of the storage architecture, the data center staff must create multiple tiers of storage on a single array with application-aware storage devices; high-cost, high-performance storage must share floor space with low-cost, high-capacity devices. Your SAN must share floor space with your NAS. Applications must be able to assign frequently accessed data to high-performance devices, while assigning archive data to low-cost storage. Furthermore, these assignments must be done automatically and dynamically, by policy, without operator intervention, in much the same fashion that the automatic transmission in your car pro-

gresses through its gears, from first to drive, without the requirement of the operator to employ the clutch and gearshift. Previous concepts of storage cost must be discarded. **The IT staff must change the storage paradigm, looking at cost of storage per IOPS², not just dollars per TB.**

In order to meet the demands of the beleaguered data center, Pillar has updated their line of storage products to simplify the deployment and management of the storage environment, improve the utilization of all storage, and deliver an application-aware, scalable architecture to enable the environment to grow, in terms of both performance and capacity.

Pillar Addresses Three 21st Century Information Challenges

(1) Distributed and Unanticipatable Demand

I/O happens in reaction to an event. When most events were predictably embedded in business process workflows, demand was, largely, predictable. Hierarchical Storage Management (HSM) could poll access patterns in a use case and place data appropriately. But, as VP Bob Maness puts it “that was past-tense management – based on what happened yesterday.” In many kinds of business today, opportunities and threats – and the consequent need for data – are often unplanned. Server consolidation in the past few years has rendered most I/O requests random. This is very different from the predictable patterns of access and sequential data access that were the norm when storage was directly attached to a server and served a single application.

“Pillar Data Systems’ *Axiom* does for storage what VMware did for servers,” observed Maness. It is the same concept – virtualizing the resources to get different service levels out of the same system. **Pillar Application Aware Profiles turn data placement into an optimization problem that computers solve well** – and the more different kinds of application data you throw on an Axiom Array, the better the optimization can become.

This capability makes Axiom an appropriate destination for multiple applications. When filled to the 80% utilization level that most Axiom customers see as normal, it can put a temporary halt to a data center’s need of more storage. Plus, if that more is Axiom storage, the habit of better storage consumption can continue.

¹ See the issue of *Clipper Notes* dated August 7, 2009, entitled *Reducing TCO of the Enterprise Data Center – Upgrading and Replacing Aging x86 Infrastructure*, and available at <http://www.clipper.com/research/TCG2009032.pdf>.

² I/Os per second.

(2) *Prioritization of Access*

Pillar Data Systems reacts to an I/O request, *as it happens*, to give it the CPU power, memory cache, and the right level of I/O to serve the data to the application with a guaranteed level of service. **An Axiom array can have up to 192 GB of non-volatile DRAM.** (More details on Axiom to follow.) **This is considerably more than the cache capacity found in arrays with traditional architecture. It is enough to use as an *allocatable and re-allocatable* resource.**

Information technology has a rich heritage in the use of cache, both on servers and on storage arrays. In Axiom, plentiful cache complements the use of SATA drives. With the recent inclusion of SSD drives, Axiom supports three tiers of data access, allowing both better prioritization and sensible use of resources.

CPU cycles, cache, and storage capacity (on the proper SATA stripe to optimize seek time, or on SSDs for still quicker access) are all part of serving I/O requests properly. **Pillar is perhaps unique in its ability to coordinate all three resources, dynamically as is needed, to support one of five service levels: Premium, High, Medium, Low, and Archive).** In addition, with the addition of more cache and CPU, Axiom can make Premium into “Super Premium,” when such service is needed.

(3) *Particularizing Restore*

Pillar Data Systems utilizes a proprietary chip for their RAID controller. Two controllers are in each Brick (see Exhibit one for definitions) and are not part of the data mover. This separation of duties, so to speak, underlies Axiom’s ability to *particularize* the restore process following a disk failure.

Pillar’s approach is to isolate the problem to a specific location, and solve it locally³. Since application data is spread across all of Axiom’s capacity, the impact of a bad area on an application is a small part of the LUN. Because each brick has two controllers, restoration is a parallelized effort. **Pillar’s cellular architecture, and the granular particularity of restore, mean that restore is fast and also has no wider impact on the system.**

³ This is like the approach that many data centers are using to deal with heat in the Data Center. Where once they would have relied on natural dissipation and room CRACs, now fans and ducting contain the heat from high-density servers and deal with it in a localized way.

Pillar’s Axiom Family

Pillar Data Systems has designed the Axiom family with a modular architecture to enable scalable growth in both capacity and performance, in parallel, with one eye constantly monitoring TCO. Having scalability, however, is only half of the battle. Pillar designed Axiom to be able to utilize up to 80% of the available capacity with their Application-Aware storage system. Where yesterday’s storage could only be tuned for cache and the type of interface, limiting typical utilization to less than 30%, Axiom arrays can differentiate services based upon application workload as defined by best practices and administrative settings specific to those workloads, creating multiple storage tiers. Applications that demand extremely low latency disk I/O can be assigned to solid-state disks (SSDs), with higher cost, while applications dealing with business continuity can be assigned to high capacity, low-cost devices, such as 1TB SATA drives. **This can result in utilization figures reaching 80%, while at the same time, reducing administrative workload and energy consumption by 50%, cutting TCO in half!**

Pillar Data Systems innovation permits the data center to provision data and applications dynamically, adjusting for any application that needs it, in the same way that servers based upon new quad-core CPUs dynamically can apply more, or less, CPU power to any application. This enables the data center to configure thinly provisioned arrays, as Axiom more reliably predicts performance requirements, simplifying storage management. If the data center can double storage utilization, this means that the enterprise can reduce capital expense for storage by half; with half the storage, the data center can reduce floor space and energy expense in half. Axiom can also improve the quality of service (QoS) for the end users, creating profiles to tune storage resources for cache, capacity, and queues, to meet business needs and ensure optimum utilization. This also frees up storage administrators to do more-productive work.

Pillar also provides a *guarantee* of high performance and utilization. Yes, I did say *guarantee*. Pillar will provide an agreed upon minimum utilization and performance service level for a full year. If any reconfiguration or other service is required to bring the Axiom to this level, then the service is delivered free of charge.

A Pillar Axiom system is made up of a variety of hardware components: a storage controller, called a *Slammer*, a storage enclosure, called

Exhibit 1 – Axiom Components

- **Axiom Slammer** – A fully-redundant *Storage Controller* to provide data movement, it allocates bandwidth according to the I/O priority of the workload; designed with a dual-processor architecture and automatic failover and restoration; includes redundant power supplies and up to 6GB of battery-backed cache memory; available with NAS or SAN protocol, either FC or iSCSI;
- **Axiom Brick** – A highly-available 13-drive *Storage Enclosure* for dual-ported SATA, FC, or SSD storage devices; Supports up to 13TB of raw storage, including one hot spare drive; includes two RAID controllers with 4MB of flash memory, 256MB ECC cache memory per controller;
- **Axiom Pilot** – This *Policy Controller* is the management interface for Axiom. Its simple, graphical controller and AxiomOne software enables policy-based provisioning with dynamic performance prioritization, fault management, capacity planning, and guided maintenance for the entire system.
- **A fabric mesh** connects the bricks and the slammer and supports data movement both in and out of the array and between bricks. This month, Axiom will support a 4 GB mesh.

Source: Pillar Data Systems

a *Brick*, a policy engine, called a *Pilot*, and connected via 4Gb mesh fabric. (See Exhibit 1, above, for the component details.)

Pillar has created two preconfigured models, the *Axiom 300*, for mid-sized deployments and budgets, and the *Axiom 600*, for business-class storage solutions that can scale to over 800TBs. Both of these models employ the Pillar *AxiomOne* software suite designed to enable customers to reduce the cost of managing information.

Pillar Axiom 300

The Axiom 300 brings enterprise-class storage functionality to mid-range applications, with that guarantee of high performance and utilization, along with an upgrade path designed to protect the initial investment made by the data center.

The Axiom 300 is pre-configured with one Slammer storage controller, with NAS, SAN, or iSCSI protocol, one Pilot policy controller, one Brick drive enclosure with either 500GB or 1TB SATA drives or 300GB or 450GB FC drives, for a maximum of 13TBs, and the *AxiomOne* software suite. It comes with a 3-year next-business-day warranty on hardware and software. The data center can also order up to three additional Bricks, either SATA or FC, and additional protocols to meet application requirements in a virtualized environment or to grow the Axiom 300 to meet future enterprise needs, up to 52TB. If your growth requirements expand beyond these limits, Pillar will refund the purchase price of software options when you upgrade to an Axiom 600, enabling the data center to move the Pilot, Bricks, and drives to the new system, protecting enterprise investment.

Pillar Axiom 600

The Axiom 600 is designed to satisfy the higher growth needs of enterprises with larger data centers or departmental requirements of larger companies. It is an ideal vehicle for consolidating multiple tiers of applications and data onto a single storage platform. The Axiom 600 stretches storage limits, doubling performance, regardless of the size and complexity of the data center environment. With Pillar's Application-Aware technology, the Axiom 600 can meet the highest utilization requirements, while managing the lowest TCO for the enterprise. The Axiom 600 can meet the high-performance needs of financial applications, OLTP, or any other application where speed is a requirement, as well as the high capacity requirements of any archiving environment.

Employing the same utilization characteristics of its little brother, the Axiom 300, the Axiom 600 can unify the storage architecture of any heterogeneous data center with both SAN and NAS requirements, deploying thin provisioning, reducing the amount of physical storage necessary to support any application, boosting utilization and reducing energy requirements. Coupled with Pillar's unique I/O queuing prioritization, the Axiom 600 enables the data center to differentiate classes of storage in a pool based upon application requirements that satisfy its performance priority. This enables the Axiom 600 to provide multiple QoS tiers, meeting a variety of application workloads.

It is a robust, business-class solution with support for up to 832TBs of raw capacity with multiple tiers of service. The highest premium

tier, Tier 0, consists of SSD or FC drives for the highest performance and availability. Intermediate tiers, Tier 1, 2, and 3, provide varying levels of performance and availability, while an archive tier, Tier 4, utilizes SATA drives for a cost effective archive workload. As a result, the Axiom 600 can reduce energy requirements and floor space by up to 50%, reducing infrastructure TCO and improving bottom-line profitability.

Designed with a modular architecture for flexibility, the Axiom 600 enjoys tremendous scalability for both capacity and performance, up to 64 Bricks and up to 4 Slammers, or storage controllers. Each Brick supports 13 disk devices with dual active RAID controllers to provide redundancy and ensure faster response to Slammer requests. Bricks can support SSD, SATA, or FC disk devices in a single storage pool, but the data center cannot mix devices within a Brick. The Axiom 600 can spread a LUN across multiple RAID controllers, enabling a distributed RAID 10 for improved performance.

The Axiom 600 is especially performant with an Oracle database. The Axiom writes in 1MB stripes, the largest stripe available. This enables Oracle to operate more efficiently as Oracle's Automatic Storage Manager also writes in 1MB stripes. A fortunate coincidence? Probably not, as Pillar is funded through a Larry Ellison company.

Pillar has also characterized the I/O of many of today's most popular applications and saved them as Axiom Profiles, available as a drop-down menu in Axiom's Storage Manager. It is easy to use, enabling an administrator to deploy a LUN with six simple clicks.

Pillar AxiomONE Software Suite

AxiomONE is an application-aware suite of software that provides the data center with an easy to use set of applications to manage and configure all aspects of your Axiom storage system. Consisting of the *AxiomONE Storage Services Manager*, the *Performance Manager*, and the *Capacity Planning Manager*, they provide resource management across the system, for CPU, cache, and capacity.

AxiomONE simplifies storage management with policy-based controls and simplifies deployment and management with an easy to use GUI. The data center can generate full copies of data with Axiom's *Volume Copy*, or utilize Pillar's data protection tools, such as *Snap FS* and *Snap LUN* snapshots. AxiomONE enables the data center to address potential problems before they

occur via the compilation of system health and performance statistics and enables proactive actions with the automatic display and notification of critical status changes and events.

Conclusion

Pillar's Axiom family increases the flexibility of data center storage assets, significantly reducing the cost and complexity of the enterprise storage architecture, with a single pool of storage and a single point of management. Pillar establishes multiple tiers of data within this pool to use the available storage resources more effectively, enabling the data center to control the TCO of the storage environment more effectively. Perhaps, Axiom can simplify your storage environment as well.



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 Authors

David Reine is Director, Enterprise Systems for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- *Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial “123” when you hear the automated attendant.)*

Anne MacFarland is Director of Information 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, Clipper Notes, 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.